



I n s i d e Q u i c k T i m e :

The QuickTime Technical Reference Library

QuickTime File Format

March 27, 2001



© Apple Computer, Inc. 2000



Apple Computer, Inc.

© 2000 Apple Computer, Inc.

All rights reserved.

No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, mechanical, electronic, photocopying, recording, or otherwise, without prior written permission of Apple Computer, Inc., with the following exceptions: Any person is hereby authorized to store documentation on a single computer for personal use only and to print copies of documentation for personal use provided that the documentation contains Apple's copyright notice.

The Apple logo is a trademark of Apple Computer, Inc.

Use of the "keyboard" Apple logo (Option-Shift-K) for commercial purposes without the prior written consent of Apple may constitute trademark infringement and unfair competition in violation of federal and state laws.

No licenses, express or implied, are granted with respect to any of the technology described in this book. Apple retains all intellectual property rights associated with the technology described in this book. This book is intended to assist application developers to develop applications only for Apple-labeled or Apple-licensed computers.

Every effort has been made to ensure that the information in this document is accurate. Apple is not responsible for typographical errors.

Apple Computer, Inc.

1 Infinite Loop

Cupertino, CA 95014

408-996-1010

Apple, the Apple logo, ColorSync, Macintosh, QuickDraw, and QuickTime are trademarks of Apple Computer, Inc., registered in the United States and other countries.

Simultaneously published in the United States and Canada.

Even though Apple has reviewed this manual, APPLE MAKES NO WARRANTY OR REPRESENTATION, EITHER EXPRESS OR IMPLIED, WITH RESPECT TO THIS MANUAL, ITS QUALITY, ACCURACY, MERCHANTABILITY, OR FITNESS FOR A PARTICULAR PURPOSE. AS A RESULT, THIS MANUAL IS SOLD "AS IS," AND YOU, THE PURCHASER, ARE ASSUMING THE ENTIRE RISK AS TO ITS QUALITY AND ACCURACY.

IN NO EVENT WILL APPLE BE LIABLE FOR DIRECT, INDIRECT, SPECIAL, INCIDENTAL, OR CONSEQUENTIAL DAMAGES RESULTING FROM ANY DEFECT OR INACCURACY IN THIS MANUAL, even if advised of the possibility of such damages.

THE WARRANTY AND REMEDIES SET FORTH ABOVE ARE EXCLUSIVE AND IN LIEU OF ALL OTHERS, ORAL OR WRITTEN, EXPRESS OR IMPLIED. No Apple dealer, agent, or employee is authorized to make any modification, extension, or addition to this warranty.

Some states do not allow the exclusion or limitation of implied warranties or liability for incidental or consequential damages, so the above limitation or exclusion may not apply to you. This warranty gives you specific legal rights, and you may also have other rights which vary from state to state.

Contents

	Figures, Tables, and Listings	11
Preface	About This Document	15
	Conventions Used in This Document	16
	Special Fonts	16
	Recent Changes	16
	Types of Notes	16
	Updates to This Book	17
	For More Information	17
Chapter 1	Introduction to Atoms	19
	Atoms	19
	Atom Layout	20
	Atom Structure	21
	QT Atoms	24
	QT Atom Containers	27
	Overview of the File Format	30
	Free Space Atoms	32
	Movie Data Atoms	33
	Preview Atoms	33
Chapter 2	Movie Atoms	35
	Overview of Movie Atoms	36
	The Movie Atom	38
	Movie Header Atoms	39
	Color Table Atoms	42
	User Data Atoms	43
	Print to Video (Full Screen Mode)	46
	Track Atoms	47
	Track Header Atoms	49

Clipping Atoms	52
Clipping Region Atoms	53
Track Matte Atoms	54
Compressed Matte Atoms	55
Edit Atoms	56
Edit List Atoms	57
Track Load Settings Atoms	59
Track Reference Atoms	60
Track Input Map Atoms	63
Media Atoms	67
Media Header Atoms	68
Handler Reference Atoms	70
Media Information Atoms	72
Video Media Information Atoms	73
Video Media Information Header Atoms	74
Sound Media Information Atoms	75
Sound Media Information Header Atoms	76
Base Media Information Atoms	77
Base Media Information Header Atoms	78
Base Media Info Atoms	79
Data Information Atoms	80
Data Reference Atoms	82
Sample Atoms	83
Sample Table Atoms	85
Sample Description Atoms	87
Time-to-Sample Atoms	88
Sync Sample Atoms	91
Sample-to-Chunk Atoms	93
Sample Size Atoms	96
Chunk Offset Atoms	97
Using Sample Atoms	99
Finding a Sample	100
Finding a Key Frame	100
Compressed Movie Resources	101
Allowing QuickTime to Compress the Movie Resource	101
Structure of a Compressed Movie Resource	101
Reference Movies	102
Reference Movie Atom	104

Reference Movie Descriptor Atom	105
Data Reference Atom	106
Data Rate Atom	107
CPU Speed Atom	108
Version Check Atom	109
Component Detect Atom	110
Component Description Record	110
Quality Atom	111

Chapter 3 Media Data Atom Types 113

Video Media	114
Video Sample Description	114
Video Sample Data	117
Uncompressed RGB	118
Uncompressed yuv2	118
JPEG	118
Motion-JPEG	119
Sound Media	124
Sound Sample Description	124
Variants of the Sample Description	125
VBR Audio (the Third Variant)	127
Sound Sample Data	127
Sound Codec Formats Supported	127
Uncompressed 8-Bit Sound	129
Uncompressed 16-Bit Sound	129
Formats Not Currently in Use: MACE 3:1 and 6:1	129
IMA, uLaw, and aLaw	129
Floating-Point Formats	130
24- and 32-Bit Integer Formats	130
kMicrosoftADPCMFormat and kDVIIIntelIMAFormat Sound	
Codecs	130
kDVAudioFormat Sound Codec	131
kQDesignCompression Sound Codec	131
MPEG Layer 3 (MP3) Codecs	131
Timecode Media	131
Timecode Sample Description	131

Timecode Media Information Atom	133
Timecode Sample Data	134
Text Media	134
Text Sample Description	134
Text Sample Data	137
Hypertext	138
Music Media	139
Music Sample Description	139
Music Sample Data	140
MPEG Media	140
MPEG Sample Description	140
MPEG Sample Data	140
Sprite Media	140
Sprite Sample Description	141
Sprite Sample Data	141
Sprite Track Properties	143
Sprite Track Media Format	145
Sprite Media Format Atoms	147
Sprite Media Format Extensions	148
Sprite Track Property Atoms	148
Atom Types	149
Sprite Button Behaviors	155
QT Atom Container Description Key	156
Sprite Media Handler Track Properties QT Atom Container Format	157
Sprite Media Handler Sample QT Atom Container Formats	157
Wired Action Grammar	159
Flash Media	167
Tween Media	168
Tween Sample Description	168
Tween Sample Data	168
Tween Type Categories	171
Tween QT Atom Container	172
General Tween Atoms	172
Path Tween Atoms	175
List Tween Atoms	176
3D Tween Atoms	176
Interpolation Tween Atoms	177
Region Tween Atoms	178

Sequence Tween Atoms	179
Modifier Tracks	180
Limitations of Spatial Modifier Tracks	181
Track References	181
Chapter Lists	182
3D Media	183
3D Sample Description	183
3D Sample Data	183
Streaming Media	184
Streaming Media Sample Description	184
Hint Media	185
Adding Hint Tracks to a Movie	186
Packetization Hint Media Header Atom	187
Hint Track User Data Atom	188
Movie Hint Info Atom	188
Finding an Original Media Track From a Hint Track	190
RTP Hint Tracks	190
Hint Sample Data Format	191
Packetization Hint Sample Data for Data Format 'rtp '	194
No-Op Data Mode	199
Immediate Data Mode	199
Sample Mode	200
Sample Description Mode	202
VR Media	203
VR World Atom Container	205
VR World Header Atom Structure	206
Imaging Parent Atom	207
Panorama-Imaging Atom	207
Node Parent Atom	209
Node Location Atom Structure	209
Custom Cursor Atoms	210
Node Information Atom Container	211
Node Header Atom Structure	212
Hot Spot Parent Atom	213
Hot Spot Information Atom	214
Specific Information Atoms	216
Link Hot Spot Atom	216
Link Hot Spot Valid Flags	217

URL Hot Spot Atom	218
Support for Wired Actions	218
QuickTime VR File Format	220
Single-Node Panoramic Movies	221
Single-Node Object Movies	222
Multinode Movies	223
QTVR Track	224
QuickTime VR Sample Description Structure	224
Panorama Tracks	225
Panorama Sample Atom Structure	225
Panorama Image Track	229
Cylindrical Panoramas	232
Cubic Panoramas	232
Image Tracks in Cubic Nodes	233
Panorama Tracks in Cubic Nodes	233
Nonstandard Cubes	235
Hot Spot Image Tracks	236
Low-Resolution Image Tracks	236
Track Reference Entry Structure	237
Object Tracks	237
Object Sample Atom Structure	238
Track References for Object Tracks	244
Movie Media	246
Movie Sample Description	246
Movie Media Sample Format	246

Chapter 4 Basic Data Types 253

Language Code Values	253
Calendar Date and Time Values	256
Matrices	256
Graphics Modes	257
RGB Colors	259
Balance	259

Chapter 5 Some Useful Examples and Scenarios 261

Creating, Copying, and Disposing of Atom Containers	262
Creating New Atoms	263
Copying Existing Atoms	265
Retrieving Atoms From an Atom Container	268
Modifying Atoms	270
Removing Atoms From an Atom Container	271
Creating an Effect Description	272
Structure of an Effect Description	273
Required Atoms of an Effects Description	274
Parameter Atoms of an Effects Description	274
Creating an Input Map	277
Structure of an Input Map	277
Building Input Maps	278
Creating Movies With Modifier Tracks	280
Authoring Movies With External Movie Targets	281
Target Atoms for Embedded Movies	282
Adding Wired Actions To a Flash Track	284
Extending the SWF Format	284
What You Need to Modify	285
File Length	285
ActionRecordsOffset	285
ActionOffset	286
Condition	286
Actions	286
DoAction	286
Creating Video Tracks at 30 Frames-per-Second	287
Creating Video Tracks at 29.97 Frames-per-Second	287
Creating Audio Tracks at 44.1 Khz	288
Creating a Timecode Track for 29.97 FPS Video	289
Playing With Edit Lists	292
Interleaving Movie Data	294
Referencing Two Data Files With a Single Track	296
Getting the Name of a QuickTime VR Node	298
Adding Custom Atoms in a QuickTime VR Movie	299
Adding Atom Containers in a QuickTime VR Movie	301
Optimizing QuickTime VR Movies for Web Playback	302

The QTVR Flattener	302
Sample Atom Container for the QTVR Flattener	304
Atom Types in QuickTime Image Files	307
Recommended File Type and Suffix	310
Digital Video File Formats	316
Digital Audio File Formats	316
Still Image File Formats	318
Animation and 3D File Formats	320

Figures, Tables, and Listings

11

Figure 1-1	A sample QuickTime atom	21
Figure 1-2	Calculating atom sizes	24
Figure 1-3	QT atom layout	26
Figure 1-4	QT atom container with parent and child atoms	28
Figure 1-5	A QT atom container with two child atoms	29
Figure 1-6	The structure of a QuickTime file	31
Table 1-1	Basic atom types of a QuickTime file	32
Figure 1-7	The layout of a preview atom	33
Figure 2-1	Sample organization of a one-track video movie	37
Figure 2-2	The layout of a movie atom	38
Figure 2-3	The layout of a movie header atom	40
Figure 2-4	The layout of a color table atom	42
Figure 2-5	The layout of a user data atom	44
Table 2-1	User data list entry types	45
Figure 2-6	The layout of a track atom	48
Figure 2-7	The layout of a track header atom	50
Figure 2-8	The layout of a clipping atom	53
Figure 2-9	The layout of a track matte atom	55
Figure 2-10	The layout of an edit atom	57
Figure 2-11	The layout of an edit list table entry	58
Figure 2-12	The layout of a track load settings atom	59
Figure 2-13	The layout of a track reference atom	61
Table 2-2	Track reference types	62
Figure 2-14	The layout of a track input map atom	64
Figure 2-15	The layout of a media atom	67
Figure 2-16	The layout of a media header atom	69
Figure 2-17	The layout of a handler reference atom	71
Figure 2-18	The layout of a media information atom for video	73
Figure 2-19	The layout of a media information header atom for video	74
Figure 2-20	The layout of a media information atom for sound	75
Figure 2-21	The layout of a sound media information header atom	76
Figure 2-22	The layout of a base media information atom	78
Figure 2-23	The layout of a base media info atom	79
Figure 2-24	The layout of a data information atom	81
Table 2-4	Data reference types	83
Figure 2-25	Samples in a media	84

Figure 2-26	The layout of a sample table atom	86
Figure 2-27	The layout of a sample description atom	87
Figure 2-28	The layout of a time-to-sample atom	89
Figure 2-29	The layout of a time-to-sample table entry	90
Figure 2-30	An example of a time-to-sample table	91
Figure 2-31	The layout of a sync sample atom	92
Figure 2-32	The layout of a sync sample table	93
Figure 2-33	The layout of a sample-to-chunk atom	94
Figure 2-34	The layout of a sample-to-chunk table entry	95
Figure 2-35	An example of a sample-to-chunk table	95
Figure 2-36	The layout of a sample size atom	96
Figure 2-37	An example of a sample size table	97
Figure 2-38	The layout of a chunk offset atom	98
Figure 2-39	An example of a chunk offset table	99
Table 2-5	Contents of complete compressed movie	102
Figure 2-40	A movie atom containing a 'rmra' atom instead of a 'mvhd' atom	103
Figure 2-41	A 'rmra' atom with multiple 'rmda' atoms	105
Figure 2-42	Reference movie descriptor atom	106
Table 3-1	Some image compression formats	115
Table 3-2	Video sample description extensions	117
Figure 3-1	Motion-JPEG A dual-field sample data	122
Figure 3-2	Motion-JPEG B dual-field sample data	123
Table 3-3	Supported QuickTime audio formats. Note that not all of these are uncompressed.	128
Table 3-4	Text sample extensions	137
Table 3-5	Sprite properties	142
Table 3-6	Sprite track properties	144
Figure 3-3	A key frame sample atom container	145
Figure 3-4	Atoms that describe a sprite and its properties	146
Figure 3-5	Atoms that describe sprite images	147
Table 3-7	Tween type values	169
Table 3-8	The 'hinf' atom type containing child atoms	189
Table 3-9	Hint track sample description	192
Table 3-10	The structure of table entries	194
Figure 3-6	Structure of the VR world atom container	206
Figure 3-7	Structure of the node information atom container	212
Figure 3-8	The structure of a single-node panoramic movie file	221
Figure 3-9	The structure of a single-node object movie file	222
Figure 3-10	The structure of a multinode movie file	224
Figure 3-11	Creating an image track for a panorama	230

Figure 3-12	Creating an image track for a panorama, with the image track oriented horizontally	231
Table 3-11	Fields and their special values as represented in the pano sample data atom, providing backward compatibility to QuickTime VR 2.2	234
Table 3-12	Values for min and max fields	235
Table 3-13	Values used for representing six square sides	236
Figure 3-13	The structure of an image track for an object	245
Table 4-1	QuickTime language code values	253
Figure 4-1	How display matrices are used in QuickTime	256
Figure 4-2	Applying the transform	257
Table 4-2	QuickTime graphics modes	258
Figure 5-1	QT atom container after inserting an atom	264
Figure 5-2	QT atom container after inserting a second atom	264
Figure 5-3	Two QT atom containers, A and B	266
Figure 5-4	QT atom container after child atoms have been inserted	267
Figure 5-5	An example effect description for the Push effect	276
Figure 5-6	An example input map referencing two sources	278
Figure 5-7	Non-interleaved movie data	295
Figure 5-8	Interleaved movie data	295
Figure A-1	An 'idsc' atom followed by an 'idat' atom	308
Table A-1	A QuickTime image file containing JPEG-compressed data	309

About This Document

This document, which supersedes the previous *QuickTime File Format Specification, June, 2000*, describes the format and content of QuickTime files. It is current as of March, 2001, and is intended primarily for developers who need to work with QuickTime files outside the context of the QuickTime environment. For example, if you are developing a non-QuickTime application that imports QuickTime files or works with QuickTime VR, you need to understand the material in this book.

QuickTime itself provides a number of high-level functions that you can use to create and manipulate QuickTime files, without requiring you to understand the actual file format. These functions serve to insulate developers from the low-level details of operation. But not all kinds of QuickTime files can be created without the information presented here.

This document begins with an introduction to QuickTime atoms, then presents the structure of the QuickTime file format in detail. This is followed by a series of code examples for manipulating a QuickTime file using the QuickTime API. A series of appendixes describes some common file formats that can be contained within a QuickTime file as data.

This document assumes that the reader is familiar with the basic concepts of digital video and digital audio, as well as with QuickTime.

This book describes QuickTime files in general, rather than how they are supported on a specific computing platform or in a specific programming language. As a result, the file format information is presented in a tabular manner, rather than in coded data structures. Similarly, field names are presented in English rather than as programming language tags. Furthermore, to the extent possible, data types are described generically. For example, this book uses “32-bit signed integer” rather than “long” to define a 32-bit integer value. Based on the information provided here, you should be able to create appropriate data structure specifications for your environment.

QuickTime files are used to store QuickTime movies, as well as other time-based data. If you are writing an application that parses QuickTime files, you should recognize that there may be non-movie data in the files.

IMPORTANT

The QuickTime file format has been used as the basis of the MPEG-4 standard, developed by the International Organization for Standardization (ISO). ▲

Conventions Used in This Document

This document uses special conventions to present certain types of information.

Special Fonts

All code listings, reserved words, and the names of actual data structures, constants, fields, parameters, and routines are shown in Letter Gothic (this is Letter Gothic).

Words that appear in **boldface** are key terms or concepts and are defined in the glossary.

Recent Changes

Additions and changes made after June, 2000 are marked with a vertical bar in the left margin, like the one on the left.

Types of Notes

There are three types of notes used in this book.

Note

A note like this contains information that is interesting but possibly not essential to an understanding of the main text. ◆

IMPORTANT

A note like this contains information that is essential for an understanding of the main text. ▲

▲ **WARNING**

A warning like this indicates potential problems that you should be aware of as you design your application. Failure to heed these warnings could result in system crashes or loss of data. ▲

Updates to This Book

For any online updates to this book, check the QuickTime developers' page on the World Wide Web at

<http://developer.apple.com/quicktime/>

or you can go directly to the QuickTime New Documentation page at

<http://developer.apple.com/techpubs/quicktime/qtdevdocs/RM/newsframe.htm>

For More Information

For information about membership in Apple's developer program, you should go to this URL:

<http://developer.apple.com/membership/>

For technical support, go to this URL:

<http://developer.apple.com/products/techsupport/>

For information on registering signatures, file types, and other technical information, contact

Macintosh Developer Technical Support
Apple Computer, Inc.
1 Infinite Loop, M/S 303-2T
Cupertino, CA 95014

P R E F A C E

Introduction to Atoms

This chapter describes how QuickTime movies are stored on disk. The QuickTime file format is designed to accommodate the various kinds of data that need to be stored in order to work with digital media. Because the file format can be used to describe almost any media structure, it is an ideal format for the exchange of digital media between applications, regardless of the platform on which the application may be running.

A QuickTime file stores the description of the media separately from the media data. The description, or meta-data, is called the **movie** and contains information such as the number of tracks, the video compression format, and timing information. The movie also contains an index of where all the media data is stored. The media data is all of the actual sample data, such as video frames and audio samples. The media data may be stored in the same file as the QuickTime movie, in a separate file, or in several files.

Before explaining the specifics of how a QuickTime movie is stored, it is important to first understand the basic units that are used to construct QuickTime files. QuickTime uses two basic structures for storing information: **classic atoms** and **QT atoms**. Both classic atoms, which are simple atoms, and QT atoms, which are atom container atoms, allow you to construct arbitrarily complex hierarchical data structures. Both also allow applications to ignore data they don't understand.

QuickTime **atom containers** provide a basic structure for storing information in QuickTime. An atom container is a tree-structured hierarchy of QT atoms.

Atoms

The basic data unit in a QuickTime file is the **atom**. Each atom contains size and type information along with its data. The size field indicates the number of

bytes in the atom, including the size and type fields. The type field specifies the type of data stored in the atom and, by implication, the format of that data.

Atom types are specified by a 32-bit integer, typically a four-character code. Apple Computer reserves all four-character codes consisting entirely of lowercase letters. Unless otherwise stated, all data in a QuickTime movie is stored in big-endian (network) byte ordering. All version fields must be set to 0, unless this document states otherwise.


Atoms are hierarchical in nature. That is, one atom can contain one or more other atoms of varying types. For example, a movie atom contains one track atom for each track in the movie. The track atoms, in turn, contain one or more media atoms each, along with other atoms that define other track and movie characteristics. This hierarchical structure of atoms is referred to as a **containment hierarchy**.

Classic atoms and QT atoms coexist within a single tree. Apple provides APIs that allow you to move within a tree of QT atoms. With classic atoms, you read their size bytes and access their contents by calculating offsets.

The format of the data stored within a given atom cannot be determined based only on the type field of that atom. That is, an atom's use is determined by its context. A given atom type can have different usages when stored within atoms of different types. This means that all QuickTime file readers must take into consideration not only the atom type, but also the atom's containment hierarchy.

Atom Layout

Figure 1-1 shows the layout of a sample QuickTime atom. Each atom carries its own size and type information as well as its data. Throughout this document, the name of a **container atom** (an atom that contains other atoms, including other container atoms) is printed in a gray box, and the name of a **leaf atom** (an atom that contains no other atoms) is printed in a white box. Leaf atoms contain data, usually in the form of tables.

Figure 1-1 A sample QuickTime atom

A leaf atom, as shown in Figure 1-1, simply contains a series of data fields accessible by offsets. You can use QuickTime's atom tools to search through QT atom hierarchies until you get to leaf atoms, then read the leaf atom's data from its various fields.

Atoms within container atoms do not have to be in any particular order, with the exception of handler description atoms. Handler description atoms must come before their data. For example, the media handler description must come before a media information atom. A data handler description atom must come before a data information atom.

Atom Structure

Atoms consist of a header, followed by atom data. The header contains the atom's size and type fields, giving the size of the atom in bytes and its type. It may also contain an extended size field, giving the size of a large atom as a 64-bit integer.

An atom header consists of the following fields.

Field descriptions

Atom size A 32-bit integer that indicates the size of the atom, including both the atom header and the atom's contents, including any contained atoms. Normally, the `size` field contains the actual size of the atom, in bytes, expressed as a 32-bit unsigned integer. However, the `size` field can contain special values that indicate an alternate method of determining the atom size. (These special values are normally used only for media data ('mdat') atoms.)

If the `size` field is set to 0, which is allowed only for a top-level atom, this is the last atom in the file and it extends to the end of the file.

If the `size` field is set to 1, then the actual size is given in the `extended size` field, an optional 64-bit field that follows the `type` field. This accommodates media data atoms that contain more than 2^{32} bytes.

Figure 1-2 shows how to calculate the size of an atom.

Type A 32-bit integer that contains the type of the atom. This can often be usefully treated as a four-character field with a mnemonic value, such as 'moov' (0x6D6F6F76) for a movie atom, or 'trak' (0x7472616B) for a track atom, but non-ASCII values (such as 0x00000001) are also used.

Knowing an atom's type allows you to interpret its data. An atom's data can be arranged as any arbitrary collection of fields, tables, or other atoms. The data structure is specific to the atom type. An atom of a given type has a defined data structure.

If your application encounters an atom of an unknown type, it should not attempt to interpret the atom's data. Use the atom's `size` field to skip this atom and all of its contents. This allows a degree of forward compatibility with extensions to the QuickTime file format.

Extended Size If the `size` field of an atom is set to 1, the `type` field is followed by a 64-bit `extended size` field, which contains the

actual size of the atom as a 64-bit unsigned integer. This is used when the size of a media data atom exceeds 2^{32} bytes.


When the `size` field contains the actual size of the atom, the `extended size` field is not present. This means that when a QuickTime atom is modified by adding data, and its size crosses the 2^{32} byte limit, there is no `extended size` field in which to record the new atom size. Consequently, it is not always possible to enlarge an atom beyond 2^{32} bytes without copying its contents to a new atom.

To prevent this inconvenience, media data atoms are typically created with a 64-bit placeholder atom immediately preceding them in the movie file. The placeholder atom has a type of `kWideAtomPlaceholderType` ('wide'). Much like a 'free' or 'skip' atom, the 'wide' atom is reserved space, but in this case the space is reserved for a specific purpose. If a 'wide' atom immediately precedes a second atom, the second atom can be extended from a 32-bit size to a 64-bit size simply by starting the atom header 8 bytes earlier (overwriting the 'wide' atom), setting the `size` field to 1, and adding an `extended size` field. This way the offsets for sample data do not need to be recalculated.

The 'wide' atom is exactly 8 bytes in size, and consists solely of its `size` and `type` fields. It contains no other data.

Note

A common error is thinking that the 'wide' atom contains the extended size. The 'wide' atom is merely a placeholder that can be overwritten if necessary, by an atom header containing an `extended size` field.

Figure 1-2 Calculating atom sizes

QT Atoms

Because of the limitations of the classic atom structure—that is, you need to know the offsets in order to move through the atom tree—Apple created QT atoms. QT atoms are an enhanced data structure that provide a more general-purpose storage format and remove some of the ambiguities that arise when using simple atoms.

In particular, with classic atoms there is no way to know if an atom contains data or whether it contains other atoms, or both, without specific knowledge about the atom. Using QT atoms, a given atom is either a leaf atom or a container atom. There is no ambiguity. Furthermore, QT atoms allow for multiple atoms of a given type to be specified through identification numbers. While QT atoms are a more powerful data structure, they require more overhead in the file.

There are some advantages to using QT atoms for holding and passing information:


- QT atoms can nest indefinitely, forming hierarchies that are easy to pass from one process to another.
- QuickTime provides a single set of tools by which you can search and manipulate QT atoms of all types.

The QuickTime file format uses both classic atoms and QT atoms. In general, newer parts of the QuickTime file format use QT atoms, while older parts use classic atoms. When defining new QuickTime structures, you should use QT atoms whenever practical.

Figure 1-3 depicts the layout of a QT atom. Each QT atom starts with a QT atom container header, followed by the root atom. The root atom's type is determined by the QT atom's type. The root atom contains any other atoms that are part of the structure.

Each container atom starts with a QT atom header followed by the atom's contents. The contents are either child atoms or data, but never both. If an atom contains children, it also contains all of its children's data and descendents. The root atom is always present and never has any siblings.

Figure 1-3 QT atom layout



A QT atom container header contains the following data.

Field descriptions

Reserved A 10-byte element that must be set to 0.
Lock count A 16-bit integer that must be set to 0.

Each QT atom header contains the following data.

Field descriptions


Size A 32-bit integer that indicates the size of the atom in bytes,
including both the QT atom header and the atom's

	contents. If the atom is a leaf atom, then this field contains the size of the single atom. The size of container atoms includes all of the contained atoms. You can walk the atom tree using the size and child count fields.
Type	A 32-bit integer that contains the type of the atom. If this is the root atom, the type value is set to 'sean'.
Atom ID	A 32-bit integer that contains the atom's ID value. This value must be unique among its siblings. The root atom always has an atom ID value of 1.
Reserved	A 16-bit integer that must be set to 0.
Child count	A 16-bit integer that specifies the number of child atoms that an atom contains. This count only includes immediate children. If this field is set to 0, the atom is a leaf atom and only contains data.
Reserved	A 32-bit integer that must be set to 0.

QT Atom Containers

A QuickTime atom container is a basic structure for storing information in QuickTime. An atom container is a tree-structured hierarchy of QT atoms. You can think of a newly created QT atom container as the root of a tree structure that contains no children.

A QT atom container contains QT atoms, as shown in Figure 1-4. Each QT atom contains either data or other atoms. If a QT atom contains other atoms, it is a **parent atom** and the atoms it contains are its **child atoms**. Each parent's child atom is uniquely identified by its **atom type** and **atom ID**. A QT atom that contains data is called a leaf atom.

Figure 1-4 QT atom container with parent and child atoms

Each QT atom has an offset that describes the atom's position within the QT atom container. In addition, each QT atom has a type and an ID. The atom type describes the kind of information the atom represents. The atom ID is used to differentiate child atoms of the same type with the same parent; an atom's ID must be unique for a given parent and type. In addition to the atom ID, each atom has a 1-based index that describes its order relative to other child atoms of the same parent with the same atom type. You can uniquely identify a QT atom in one of three ways:

- by its offset within its QT atom container
- by its parent atom, type, and index
- by its parent atom, type, and ID

You can store and retrieve atoms in a QT atom container by index, ID, or both. For example, to use a QT atom container as a dynamic array or tree structure, you can store and retrieve atoms by index. To use a QT atom container as a database, you can store and retrieve atoms by ID. You can also create, store, and retrieve atoms using both ID and index to create an arbitrarily complex, extensible data structure.

▲ **WARNING**


Since QT atoms are offsets into a data structure, they can be changed during editing operations on QT atom containers, such as inserting or deleting atoms. For a given atom, editing child atoms is safe, but editing sibling or parent atoms invalidates that atom's offset. ▲

Note

For cross-platform purposes, all data in a QT atom is expected to be in big-endian format. However, leaf data can be little-endian if it is custom to an application. ♦

Figure 1-5 shows a QT atom container that has two child atoms. The first child atom (offset = 10) is a leaf atom that has an atom type of 'abcd', an ID of 1000, and an index of 1. The second child atom (offset = 20) has an atom type of 'abcd', an ID of 900, and an index of 2. Because the two child atoms have the same type, they must have different IDs. The second child atom is also a parent atom of three atoms.

Figure 1-5 A QT atom container with two child atoms



The first child atom (offset = 30) has an atom type of 'abcd', an ID of 100, and an index of 1. It does not have any children, nor does it have data. The second child atom (offset = 40) has an atom type of 'word', an ID of 100, and an index of 1. The atom has data, so it is a leaf atom. The second atom (offset = 40) has the same ID as the first atom (offset = 30), but a different atom type. The third child atom (offset = 50) has an atom type of 'abcd', an ID of 1000, and an index of 2. Its atom type and ID are the same as that of another atom (offset = 10) with a different parent.

Note

If you are working with the QuickTime API, you do not need to parse QT atoms. Instead, the QT atom functions can be used to create atom containers, add atoms to and remove atoms from atom containers, search for atoms in atom containers, and retrieve data from atoms in atom containers. ♦


Most QT atom functions take two parameters to specify a particular atom: the atom container that contains the atom, and the offset of the atom in the atom container data structure. You obtain an atom's offset by calling either `QTFindChildByID` or `QTFindChildByIndex`. An atom's offset may be invalidated if the QT atom container that contains it is modified.

When calling any QT atom function for which you specify a parent atom as a parameter, you can pass the constant `kParentAtomIsContainer` as an atom offset to indicate that the specified parent atom is the atom container itself. For example, you would call the `QTFindChildByIndex` function and pass `kParentAtomIsContainer` constant for the parent atom parameter to indicate that the requested child atom is a child of the atom container itself.

Overview of the File Format

A QuickTime file is simply a collection of atoms. QuickTime does not impose any rules about the order of these atoms.

Figure 1-6 depicts a typical QuickTime file.

Figure 1-6 The structure of a QuickTime file

In file systems that support filename extensions, QuickTime filenames typically have an extension of `.mov`. On the Macintosh platform, QuickTime files have a file type of `'Moov'`. On the Macintosh, the movie atom may be stored as a Macintosh resource using the Resource Manager. The resource has a type of `'moov'`. All media data is stored in the data fork. On the Internet, QuickTime files are served under the mime type `"video/quicktime"`.

As previously discussed, QuickTime files consist of atoms, each with an appropriate atom type. A few of these types are considered basic atom types

and form the structure within which the other atoms are stored. Table 1-1 lists the currently supported basic atom types.

Table 1-1 Basic atom types of a QuickTime file

Atom type	Use
'free'	Unused space available in file.
'skip'	Unused space available in file.
'wide'	Reserved space—can be overwritten by an extended size field.
'pnot'	Reference to movie preview data.
'moov'	Movie resource—meta-data about the movie (number and type of tracks, location of sample data, and so on).
'mdat'	Movie sample data—usually this data can be interpreted only by using the movie resource.

Although it is true that QuickTime imposes no strict order on a movie's atoms, it is often convenient if the movie atom appears near the front of the file. For example, an application that plays a movie over a network would not necessarily have access to the entire movie at all times. If the movie atom is stored at the beginning of the file, the application can use the meta-data to understand the movie's content as it is acquired over the network. This can allow an application to play a movie as it downloads.

The following sections, as well as Chapter 2, "Movie Atoms," describe each of these basic atom types in more detail, including descriptions of the atoms that each basic atom may contain.

Free Space Atoms

Both free and skip atoms designate unused space in the movie data file. These atoms consist only of an atom header (atom size and type fields), followed by the appropriate number of bytes of free space. When reading a QuickTime movie, your application may safely skip these atoms. When writing or updating a movie, you may reuse the space associated with these atom types. The wide atom consists only of a type and size field. This occupies 8 bytes—enough space to add an extended size field to the header of the atom that follows. If an atom

grows to exceed 2^{32} bytes in size, and it is preceded by a wide atom, you may create a new atom header, containing an extended size field, by overwriting the wide atom.

Movie Data Atoms

As with the free and skip atoms, the movie data atom is structured quite simply. It consists of an atom header (atom size and type fields), followed by the movie's media data. Your application can understand the data in this atom only by using the meta-data stored in the movie atom. This atom can be quite large, and may exceed 2^{32} bytes, in which case the size field will be set to 1, and the header will contain a 64-bit extended size field.

Preview Atoms

The preview atom contains information that allows you to find the preview image associated with a QuickTime movie. The preview image, or poster, is a representative image suitable for display to the user in, for example, Open dialog boxes. Figure 1-7 depicts the layout of the preview atom.

Figure 1-7 The layout of a preview atom

Bytes	
Preview atom	
Atom size	4
Type = 'pnot'	4
Modification date	4
Version number	2
Atom type	4
Atom index	2

The preview atom has an atom type value of 'pnot' and, following its atom header, contains the following fields.

Field descriptions

Size	A 32-bit integer that specifies the number of bytes in this preview atom.
Type	A 32-bit integer that identifies the atom type; this field must be set to 'pnot'.
Modification date	A 32-bit unsigned integer containing a date that indicates when the preview was last updated. The data is in standard Macintosh format.
Version number	A 16-bit integer that must be set to 0.
Atom type	A 32-bit integer that indicates the type of atom that contains the preview data. Typically, this is set to 'PICT' to indicate a QuickDraw picture.
Atom index	A 16-bit integer that identifies which atom of the specified type is to be used as the preview. Typically, this field is set to 1 to indicate that you should use the first atom of the type specified in the atom type field.

Movie Atoms

This chapter provides a general introduction to QuickTime movie atoms, as well as specific details on the layout and usage of these atoms.

If you are a QuickTime application or tool developer, you'll want to read this chapter in order to understand the characteristics and usage of QuickTime movie atoms. Each atom type discussed in this chapter is shown with an accompanying illustration that contains offset information, followed by field descriptions.

This chapter is divided into the following major sections:

- “Overview of Movie Atoms” (page 36) discusses QuickTime movie atoms, which act as containers for information that describes a movie's data. A conceptual illustration is provided that shows the organization of a simple, one-track QuickTime movie. Color table atoms and user data atoms are also discussed.
- “Track Atoms” (page 47) describes track atoms, which define a single track of a movie.
- “Media Atoms” (page 67) discusses media atoms, which define a track's movie data.
- “Sample Atoms” (page 83) discusses how QuickTime stores media data in samples. The section also includes examples of how you use these sample atoms.
- “Compressed Movie Resources” (page 101) discusses compressed movie resources, which are made up of a group of QuickTime atoms arranged in a hierarchy.
- “Reference Movies” (page 102) discusses movies that contain a list of references to alternate movies, as well as the criteria for selecting the correct movie from a list of alternates.


Overview of Movie Atoms

QuickTime movie atoms have an atom type of 'moov'. These atoms act as a container for the information that describes a movie's data. This information, or meta-data, is stored in a number of different types of atoms.

The movie atom is essentially a container of other atoms. At the highest level, movie atoms generally contain track atoms, which in turn contain media atoms. At the lowest level you find the leaf atoms, which contain actual data, usually in the form of a table or a data stream. For example, the track atom contains the edit atom, which contains a leaf atom called the *edit list atom*. The edit list atom contains an edit list table. Both of these atoms are discussed later in this book.

Figure 2-1 provides a conceptual view of the organization of a simple, one-track QuickTime movie. Each nested box in the illustration represents an atom that belongs to its parent atom. The figure does not show the data regions of any of the atoms. These areas are described in the sections that follow.

Figure 2-1 Sample organization of a one-track video movie




The Movie Atom

You use movie atoms to specify the information that defines a movie—that is, the information that allows your application to understand the data that is stored in the movie data atom. The movie atom normally contains a movie header atom, which defines the time scale and duration information for the entire movie, as well as its display characteristics. In addition, the movie atom contains each track in the movie.

The movie atom has an atom type of 'moov'. It contains other types of atoms, including at least one parent atom—either the movie header atom ('mvhd'), or a reference movie atom ('rmra'). A movie atom can contain both a movie header atom and a reference movie atom, but it must contain at least one of the two. It can also contain several other atoms, such as a clipping atom ('clip'), one or more track atoms ('trak'), a color table atom ('ctab'), and a user data atom ('udta').

Figure 2-2 shows the layout of a typical movie atom.

Figure 2-2 The layout of a movie atom



Required atom

A movie atom may contain the following information.

Movie Atoms

Field descriptions


Size	The number of bytes in this movie atom.
Type	The type of this movie atom; this field must be set to 'moov'.
Movie header atom	See “Movie Header Atoms” (page 39) for more information.
Movie clipping atom	See “Clipping Atoms” (page 52) for more information.
Track list atoms	See “Track Atoms” (page 47) for details on track atoms and their associated atoms.
User data atom	See “User Data Atoms” (page 43) for more information about user data atoms.
Color table atom	See “Color Table Atoms” (page 42) for a discussion of the color table atom.
Reference movie atom	See “Reference Movies” (page 102) for a discussion of reference movie atoms.

Movie Header Atoms

You use the movie header atom to specify the characteristics of an entire QuickTime movie. The data contained in this atom defines characteristics of the entire QuickTime movie, such as time scale and duration. It has an atom type value of 'mvhd'.

Figure 2-3 shows the layout of the movie header atom. The movie header atom is a leaf atom.

Figure 2-3 The layout of a movie header atom



You define a movie header atom by specifying the following data elements.

Field descriptions

Size	A 32-bit integer that specifies the number of bytes in this movie header atom.
Type	A 32-bit integer that identifies the atom type; must be set to 'mvhd'.
Version	A 1-byte specification of the version of this movie header atom.
Flags	Three bytes of space for future movie header flags.
Creation time	A 32-bit integer that specifies (in seconds since midnight, January 1, 1904) when the movie atom was created.
Modification time	A 32-bit integer that specifies (in seconds since midnight, January 1, 1904) when the movie atom was changed.
Time scale	A time value that indicates the time scale for this movie—that is, the number of time units that pass per second in its time coordinate system. A time coordinate system that measures time in sixtieths of a second, for example, has a time scale of 60.
Duration	A time value that indicates the duration of the movie in time scale units. Note that this property is derived from the movie's tracks. The value of this field corresponds to the duration of the longest track in the movie.
Preferred rate	A 32-bit fixed-point number that specifies the rate at which to play this movie. A value of 1.0 indicates normal rate.
Preferred volume	A 16-bit fixed-point number that specifies how loud to play this movie's sound. A value of 1.0 indicates full volume.
Reserved	Ten bytes reserved for use by Apple. Set to 0.
Matrix structure	The matrix structure associated with this movie. A matrix shows how to map points from one coordinate space into another. See “Matrices” (page 256) for a discussion of how display matrices are used in QuickTime.
Preview time	The time value in the movie at which the preview begins.
Preview duration	The duration of the movie preview in movie time scale units.
Poster time	The time value of the time of the movie poster.
Selection time	The time value for the start time of the current selection.
Selection duration	The duration of the current selection in movie time scale units.

Movie Atoms


Current time	The time value for current time position within the movie.
Next track ID	A 32-bit integer that indicates a value to use for the track ID number of the next track added to this movie. Note that 0 is not a valid track ID value.

Color Table Atoms

Color table atoms define a list of preferred colors for displaying the movie on devices that support only 256 colors. The list may contain up to 256 colors. These optional atoms have a type value of 'ctab'. The color table atom contains a Macintosh color table data structure.

Figure 2-4 shows the layout of the color table atom.

Figure 2-4 The layout of a color table atom



The color table atom contains the following data elements.

Field descriptions

Size	A 32-bit integer that specifies the number of bytes in this color table atom.
Type	A 32-bit integer that identifies the atom type; this field must be set to 'ctab'.
Color table seed	A 32-bit integer that must be set to 0.
Color table flags	A 16-bit integer that must be set to 0x8000.

Movie Atoms

Color table size	A 16-bit integer that indicates the number of colors in the following color array. This is a zero-relative value; setting this field to 0 means that there is one color in the array.
Color array	An array of colors. Each color is made of four unsigned 16-bit integers. The first integer must be set to 0, the second is the red value, the third is the green value, and the fourth is the blue value.

User Data Atoms

User data atoms allow you to define and store data associated with a QuickTime object, such as a movie, track, or media. This includes both information that QuickTime looks for, such as copyright information or whether a movie should loop, and arbitrary information—provided by and for your application—that QuickTime simply ignores.


A user data atom whose immediate parent is a movie atom contains data relevant to the movie as a whole. A user data atom whose parent is a track atom contains information relevant to that specific track. A QuickTime movie file may contain many user data atoms, but only one user data atom is allowed as the immediate child of any given movie atom or track atom.

The user data atom has an atom type of 'udta'. Inside the user data atom is a list of atoms describing each piece of user data. User data provides a simple way to extend what is stored in a QuickTime movie. For example, you can use data atoms to store a movie's window position, playback characteristics, or creation information.

This section describes the data atoms that QuickTime recognizes. You may create new data atom types that your own application recognizes. Applications should ignore any data atom types they do not recognize.

Figure 2-5 shows the layout of a user data atom.

Figure 2-5 The layout of a user data atom



The user data atom contains the following data elements.

Field descriptions

Size	A 32-bit integer that specifies the number of bytes in this user data atom.
Type	A 32-bit integer that identifies the atom type; this field must be set to 'udta'.
User data list	A user data list that is formatted as a series of atoms. Each data element in the user data list contains size and type information along with its data. For historical reasons, the data list is optionally terminated by a 32-bit integer set to 0. If you are writing a program to read user data atoms, you should allow for the terminating 0. However, if you are writing a program to create user data atoms, you can safely leave out the trailing 0.

Table 2-1 lists the currently defined list entry types.

Table 2-1 User data list entry types

List entry type	Description
'@cpy'	Copyright statement.
'@day'	Date the movie content was created.
'@dir'	Name of movie's director.
'@ed1' to '@ed9'	Edit dates and descriptions.
'@fmt'	Indication of movie format (computer-generated, digitized, and so on).
'@inf'	Information about the movie.
'@prd'	Name of movie's producer.
'@prf'	Names of performers.
'@req'	Special hardware and software requirements.
'@src'	Credits for those who provided movie source content.
'@wrt'	Name of movie's writer.
'name'	Name of object.
'WLOC'	Default window location for movie. Two 16-bit values, {x,y}.
'LOOP'	Looping. This atom is not present unless the movie is set to loop. Long integer indicating looping style. 0 for normal looping, 21 for palindromic looping.
'Sel0'	Play selection only. Byte indicating that only the selected area of the movie should be played.
'AllF'	Play all frames. Byte indicating that all frames of video should be played, regardless of timing.

Table 2-1 User data list entry types

'ptv '	Print to video. Display movie in full screen mode. This atom contains a 16-byte structure, described in “Print to Video (Full Screen Mode)” (page 46).
'hnti'	Hint info atom. Data used for real-time streaming of a movie or a track. For more information, see “Movie Hint Info Atom” (page 188) and “Hint Track User Data Atom” (page 188).
'hinf'	Hint track information. Statistical data for real-time streaming of a particular track. For more information, see “Hint Track User Data Atom” (page 188).

All user data list entries whose type begins with the © character (ASCII 169) are defined to be international text. These list entries must contain a list of text strings with associated language codes. By storing multiple versions of the same text, a single user data text item can contain translations for different languages.

The list of text strings uses a small integer atom format, which is identical to the QuickTime atom format, except that it uses 16-bit values for size and type instead of 32-bit values. The first value is the size of the string, including the size and type, and the second value is the language code for the string.

The window location, looping, play selection only, play all frames, and print to video atoms control the way QuickTime displays a movie. These atoms are interpreted *only* if the user data atom’s immediate parent is a movie atom ('moov'). If they are included as part of a track atom’s user data, they are ignored.

Print to Video (Full Screen Mode)

A movie atom’s user data atom may contain a print to video atom ('ptv '). If a print to video atom is present, QuickTime plays the movie in full-screen mode, with no window and no visible controller. Any portion of the screen not occupied by the movie is cleared to black. The user must press the `esc` key to exit full-screen mode.

This atom is often added and removed transiently to control the display mode of a movie for a single presentation, but it can also be stored as part of the permanent movie file.

The print to video atom consists of the following data.

Field descriptions

Size	A 32-bit integer that specifies the number of bytes in this user data item.
Type	A 32-bit integer that identifies the item type; this field must be set to 'ptv '. Note that the fourth character is an ASCII blank.
Display size	A 16-bit <i>little-endian</i> integer indicating the display size for the movie: 0 indicates that the movie should be played at its normal size; 1 indicates that the movie should be played at double size; 2 indicates that the movie should be played at half size; 3 indicates that the movie should be scaled to fill the screen; 4 indicates that the movie should be played at its current size (this last value is normally used when the print to video atom is inserted transiently and the movie has been temporarily resized).
Reserved1	A 16-bit integer whose value should be 0.
Reserved2	A 16-bit integer whose value should be 0.
Slide show	An 8-bit boolean whose value is 1 for a slide show. In slide show mode, the movie advances one frame each time the right-arrow key is pressed. Audio is muted.
Play on open	An 8-bit boolean whose value is normally 1, indicating that the movie should play when opened. Since there is no visible controller in full-screen mode, applications should always set this field to 1 to prevent user confusion.

Track Atoms

Track atoms define a single track of a movie. A movie may consist of one or more tracks. Each track is independent of the other tracks in the movie and carries its own temporal and spatial information. Each track atom contains its associated media atom.

Tracks are used specifically for the following purposes:

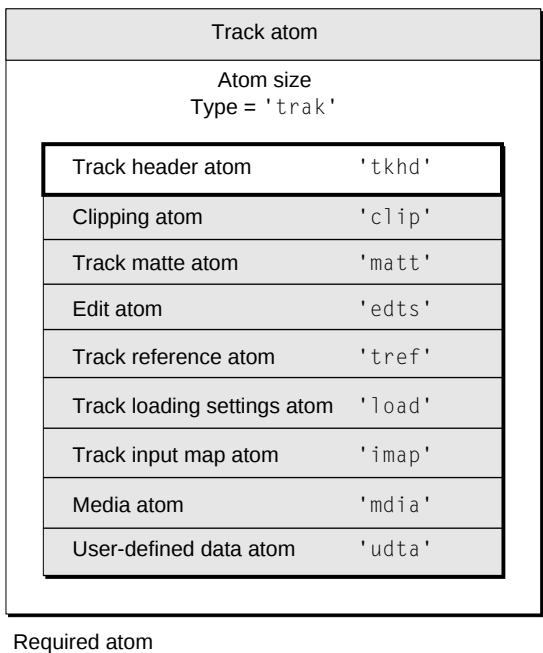
- To contain media data (media tracks).
- To contain modifier tracks (tweens, and so forth).

- To contain packetization information for streaming protocols (hint tracks).
For more information about hint tracks, refer to “Hint Media” (page 185).

Note that there must be at least one media track within a QuickTime movie. All media tracks that are present must remain in the movie, even if the media data within them is not referenced by the hint tracks. After deleting all hint tracks, the entire unhinted movie must remain.

Figure 2-6 shows the layout of a track atom. Track atoms have an atom type value of 'trak'. The track atom requires the track header atom ('tkhd') and the media atom ('mdia'). Other child atoms are optional, and may include a track clipping atom ('clip'), a track matte atom ('matt'), an edit atom ('edts'), a track reference atom ('tref'), a track load settings atom ('load'), a track input map atom ('imap'), and a user data atom ('udta').

Figure 2-6 The layout of a track atom



Track atoms contain the following data elements.

Field descriptions

Size	A 32-bit integer that specifies the number of bytes in this track atom.
Type	A 32-bit integer that identifies the atom type; this field must be set to 'trak'.
Track header atom	See “Track Header Atoms” (page 49) for details.
Clipping atom	See “Clipping Atoms” (page 52) for more information.
Track matte atom	See “Track Matte Atoms” (page 54) for more information.
Edit atom	See “Edit Atoms” (page 56) for details.
Track reference atom	See “Track Reference Atoms” (page 60)” for details.
Track load settings atom	See “Track Load Settings Atoms” (page 59) for details.
Track input map atom	See “Track Input Map Atoms” (page 63)” for details.
Media atom	See “Media Atoms” (page 67) for details.
User-defined data atom	See “User Data Atoms” (page 43) for more information.

Track Header Atoms

The track header atom specifies the characteristics of a single track within a movie. A track header atom contains a size field that specifies the number of bytes and a type field that indicates the format of the data (defined by the atom type 'tkhd').

Figure 2-7 shows the structure of the track header atom.

Figure 2-7 The layout of a track header atom

Bytes	
Track header atom	
Atom size	4
Type = 'tkhd'	4
Version	1
Flags	3
Creation time	4
Modification time	4
Track ID	4
Reserved	4
Duration	4
Reserved	8
Layer	2
Alternate group	2
Volume	2
Reserved	2
Matrix structure	36
Track width	4
Track height	4

The track header atom contains the track characteristics for the track, including temporal, spatial, and volume information.

Track header atoms contain the following data elements.

Field descriptions

Size	A 32-bit integer that specifies the number of bytes in this track header atom.
------	--

Movie Atoms

Type	A 32-bit integer that identifies the atom type; this field must be set to 'tkhd'.
Version	A 1-byte specification of the version of this track header.
Flags	Three bytes that are reserved for the track header flags. These flags indicate how the track is used in the movie. The following flags are valid (all flags are enabled when set to 1). Track enabled Indicates that the track is enabled. Flag value is 0x0001. Track in movie Indicates that the track is used in the movie. Flag value is 0x0002. Track in preview Indicates that the track is used in the movie's preview. Flag value is 0x0004. Track in poster Indicates that the track is used in the movie's poster. Flag value is 0x0008.
Creation time	A 32-bit integer that indicates (in seconds since midnight, January 1, 1904) when the track header was created.
Modification time	A 32-bit integer that indicates (in seconds since midnight, January 1, 1904) when the track header was changed.
Track ID	A 32-bit integer that uniquely identifies the track. The value 0 cannot be used.
Reserved	A 32-bit integer that is reserved for use by Apple. Set this field to 0.
Duration	A time value that indicates the duration of this track (in the movie's time coordinate system). Note that this property is derived from the track's edits. The value of this field is equal to the sum of the durations of all of the track's edits. If there is no edit list, then the duration is the sum of the sample durations, converted into the movie timescale.
Reserved	An 8-byte value that is reserved for use by Apple. Set this field to 0.

Movie Atoms


Layer	A 16-bit integer that indicates this track's spatial priority in its movie. The QuickTime Movie Toolbox uses this value to determine how tracks overlay one another. Tracks with lower layer values are displayed in front of tracks with higher layer values.
Alternate group	A 16-bit integer that specifies a collection of movie tracks that contain alternate data for one another. QuickTime chooses one track from the group to be used when the movie is played. The choice may be based on such considerations as playback quality, language, or the capabilities of the computer.
Volume	A 16-bit fixed-point value that indicates how loudly this track's sound is to be played. A value of 1.0 indicates normal volume.
Reserved	A 16-bit integer that is reserved for use by Apple. Set this field to 0.
Matrix structure	The matrix structure associated with this track. See Figure 4-1 (page 256) for an illustration of a matrix structure.
Track width	A 32-bit fixed-point number that specifies the width of this track in pixels.
Track height	A 32-bit fixed-point number that indicates the height of this track in pixels.

Clipping Atoms

Clipping atoms specify the clipping regions for movies and for tracks. The clipping atom has an atom type value of 'clip'.

Figure 2-8 shows the layout of a clipping atom.

Figure 2-8 The layout of a clipping atom



Clipping atoms contain the following data elements.

Field descriptions

- Size** A 32-bit integer that specifies the number of bytes in this clipping atom.
- Type** A 32-bit integer that identifies the atom type; this field must be set to 'clip'.
- Clipping region atom** See “Clipping Region Atoms” (page 53).

Clipping Region Atoms

The clipping region atom contains the data that specifies the clipping region, including its size, bounding box, and region. Clipping region atoms have an atom type value of 'crgn'.

The layout of the clipping region atom is shown in Figure 2-8.

Clipping region atoms contain the following data elements.

Movie Atoms

Field descriptions

Size	A 32-bit integer that specifies the number of bytes in this clipping region atom.
Type	A 32-bit integer that identifies the atom type; this field must be set to 'crgn'.
Region size	The region size, region boundary box, and clipping region data fields constitute a QuickDraw region.
Region boundary box	The region size, region boundary box, and clipping region data fields constitute a QuickDraw region.
Clipping region data	The region size, region boundary box, and clipping region data fields constitute a QuickDraw region.


Track Matte Atoms

Track matte atoms are used to visually blend the track's image when it is displayed.

Track matte atoms have an atom type value of 'matt'.

Figure 2-9 shows the layout of track matte atoms.

Figure 2-9 The layout of a track matte atom



Track matte atoms contain the following data elements.

Field descriptions

Size A 32-bit integer that specifies the number of bytes in this track matte atom.

Type A 32-bit integer that identifies the atom type; this field must be set to 'matt'.

Compressed matte atom

The actual matte data
See “Compressed Matte Atoms” (page 55) for details.

Compressed Matte Atoms

The compressed matte atom specifies the image description structure and the matte data associated with a particular matte atom. Compressed matte atoms have an atom type value of 'kmat'.

The layout of the compressed matte atom is shown in Figure 2-9.

Compressed matte atoms contain the following data elements.

Field descriptions

Size	A 32-bit integer that specifies the number of bytes in this compressed matte atom.
Type	A 32-bit integer that identifies the atom type; this field must be set to 'kmat'.
Version	A 1-byte specification of the version of this compressed matte atom.
Flags	Three bytes of space for flags. Set this field to 0.
Matte image description structure	An image description structure associated with this matte data. The image description contains detailed information that governs how the matte data is used. See “Video Sample Description” (page 114) for more information about image descriptions.
Matte data	The compressed matte data, which is of variable length.

Edit Atoms

You use edit atoms to define the portions of the media that are to be used to build up a track for a movie. The edits themselves are contained in an edit list table, which consists of time offset and duration values for each segment. Edit atoms have an atom type value of 'edts'.


Figure 2-10 shows the layout of an edit atom.

In the absence of an edit list, the presentation of a track starts immediately. An empty edit is used to offset the start time of a track.

Note

If the edit atom or the edit list atom is missing, you can assume that the entire media is used by the track. ♦

Figure 2-10 The layout of an edit atom



Edit atoms contain the following data elements.

Field descriptions

Size	A 32-bit integer that specifies the number of bytes in this edit atom.
Type	A 32-bit integer that identifies the atom type; this field must be set to 'edts'.
Edit list atom	See “Edit List Atoms” (page 57).

Edit List Atoms

You use the edit list atom, also shown in Figure 2-10, to map from a time in a movie to a time in a media, and ultimately to media data. This information is in the form of entries in an edit list table, shown in Figure 2-11. Edit list atoms have an atom type value of 'elst'.

Edit list atoms contain the following data elements.

Field descriptions

Size	A 32-bit integer that specifies the number of bytes in this edit list atom.
Type	A 32-bit integer that identifies the atom type; this field must be set to 'elst'.
Version	A 1-byte specification of the version of this edit list atom.
Flags	Three bytes of space for flags. Set this field to 0.
Number of entries	A 32-bit integer that specifies the number of entries in the edit list atom that follows.
Edit list table	An array of 32-bit values, grouped into entries containing 3 values each. Figure 2-11 shows the layout of the entries in this table.

Figure 2-11 The layout of an edit list table entry

Track duration	Media time	Media rate	Field
4	4	4	Bytes

An edit list table entry contains the following elements.

Field descriptions

Track duration	A 32-bit integer that specifies the duration of this edit segment in units of the movie's time scale.
Media time	A 32-bit integer containing the starting time within the media of this edit segment (in media timescale units). If this field is set to -1, it is an empty edit. The last edit in a track should never be an empty edit. Any difference between the movie's duration and the track's duration is expressed as an implicit empty edit.
Media rate	A 32-bit fixed-point number that specifies the relative rate at which to play the media corresponding to this edit segment. This rate value cannot be 0 or negative.

Track Load Settings Atoms

Track load settings atoms contain information that indicates how the track is to be used in its movie. Applications that read QuickTime files can use this information to process the movie data more efficiently. Track load settings atoms have an atom type value of 'load'.

Figure 2-12 The layout of a track load settings atom

Track load settings atom		Bytes
Atom size		4
Type = 'load'		4
Preload start time		4
Preload duration		4
Preload flags		4
Default hints		4

Track load settings atoms contain the following data elements.

Field descriptions

Size	A 32-bit integer that specifies the number of bytes in this track load settings atom.
Type	A 32-bit integer that identifies the atom type; this field must be set to 'load'.
Preload start time	A 32-bit integer specifying the starting time, in the movie's time coordinate system, of a segment of the track that is to be preloaded. Used in conjunction with the preload duration.
Preload duration	A 32-bit integer specifying the duration, in the movie's time coordinate system, of a segment of the track that is to be preloaded. If the duration is set to -1, it means that the preload segment extends from the preload start time to the end of the track. All media data in the segment of the track defined by the preload start time and preload duration

	values should be loaded into memory when the movie is to be played.
Preload flags	A 32-bit integer containing flags governing the preload operation. Only two flags are defined, and they are mutually exclusive. If preload flags is set to 1, the track is to be preloaded regardless of whether it is enabled. If preload flags is set to 2, the track is to be preloaded only if it is enabled.
Default hints	A 32-bit integer containing playback hints. More than one flag may be enabled. Flags are enabled by setting them to 1. The following flags are defined.
	Double buffer This flag indicates that the track should be played using double-buffered I/O. This flag's value is 0x0020. High quality This flag indicates that the track should be displayed at highest possible quality, without regard to real-time performance considerations. This flag's value is 0x0100.

Track Reference Atoms

Track reference atoms define relationships between tracks. Track reference atoms allow one track to specify how it is related to other tracks. For example, if a movie has three video tracks and three sound tracks, track references allow you to identify the related sound and video tracks. Track reference atoms have an atom type value of 'tref'.

Track references are uni-directional and point from the recipient track to the source track. For example, a video track may reference a time code track to indicate where its time code is stored, but the time code track would not reference the video track. The time code track is the source of time information for the video track.

A single track may reference multiple tracks. For example, a video track could reference a sound track to indicate that the two are synchronized and a time code track to indicate where its time code is stored.


Movie Atoms

A single track may also be referenced by multiple tracks. For example, both a sound and video track could reference the same time code track if they share the same timing information.

If this atom is not present, the track is not referencing any other track in any way. Note that the array of track reference type atoms is sized to fill the track reference atom. Track references with a reference index of 0 are permitted. This indicates no reference.

Figure 2-13 shows the layout of a track reference atom.

Figure 2-13 The layout of a track reference atom



A track reference atom contains the following data elements.

Field descriptions

- Size** A 32-bit integer that specifies the number of bytes in this track reference atom.
- Type** A 32-bit integer that identifies the atom type; this field must be set to 'tref'.

Track reference type atoms

A list of track reference type atoms containing the track reference information. These atoms are described next.

Each track reference atom defines relationships with tracks of a specific type. The reference type implies a track type. Table 2-2 shows the track reference types and their descriptions.

Table 2-2 Track reference types

Reference type	Description
'tmcd'	Time code. Usually references a time code track.
'chap'	Chapter or scene list. Usually references a text track.
'sync'	Synchronization. Usually between a video and sound track. Indicates that the two tracks are synchronized. The reference can be from either track to the other, or there may be two references.
'scpt'	Transcript. Usually references a text track.
'ssrc'	Nonprimary source. Indicates that the referenced track should send its data to this track, rather than presenting it. The referencing track will use the data to modify how it presents its data. See “Track Input Map Atoms” (page 63) for more information.
'hint'	The referenced tracks contain the original media for this hint track.

Each track reference type atom contains the following data elements.

Field descriptions

Size	A 32-bit integer that specifies the number of bytes in this track reference type atom.
Type	A 32-bit integer that identifies the atom type; this field must be set to one of the values shown in Table 2-2.
Track IDs	A list of track ID values (32-bit integers) specifying the related tracks. Note that this is one case where track ID values can be set to 0. Unused entries in the atom may have a track ID value of 0. Setting the track ID to 0 may be more convenient than deleting the reference.


You can determine the number of track references stored in a track reference type atom by subtracting its header size from its overall size and then dividing by the size, in bytes, of a track ID.

Track Input Map Atoms

Track input map atoms define how data being sent to this track from its nonprimary sources is to be interpreted. Track references of type 'ssrc' define a track's secondary data sources. These sources provide additional data that is to be used when processing the track. Track input map atoms have an atom type value of 'imap'.

Figure 2-14 shows the layout of a track input atom. This atom contains one or more track input atoms. Note that the track input map atom is a QT atom structure.

Figure 2-14 The layout of a track input map atom



Each track input map atom contains the following data elements.

Movie Atoms

Field descriptions

Size	A 32-bit integer that specifies the number of bytes in this track input map atom.
Type	A 32-bit integer that identifies the atom type; this field must be set to 'imap'.
Track input atoms	A list of track input atoms specifying how to use the input data.

The input map defines all of the track's secondary inputs. Each secondary input is defined using a separate track input atom.

Each track input atom contains the following data elements.

Field descriptions

Size	A 32-bit integer that specifies the number of bytes in this track input atom.
Type	A 32-bit integer that identifies the atom type; this field must be set to 'in' (note that the two leading bytes must be set to 0x00).
Atom ID	A 32-bit integer relating this track input atom to its secondary input. The value of this field corresponds to the index of the secondary input in the track reference atom. That is, the first secondary input corresponds to the track input atom with an atom ID value of 1; the second to the track input atom with an atom ID of 2, and so on.
Reserved	A 16-bit integer that must be set to 0.
Child count	A 16-bit integer specifying the number of child atoms in this atom.
Reserved	A 32-bit integer that must be set to 0.

The track input atom, in turn, may contain two other types of atoms: input type atoms and object ID atoms. The input type atom is required; it specifies how the data is to be interpreted.

The input type atom contains the following data elements.

Field descriptions

Size	A 32-bit integer that specifies the number of bytes in this input type atom.
------	--

Type	A 32-bit integer that identifies the atom type; this field must be set to ' ty' (note that the two leading bytes must be set to 0x00).
Input type	A 32-bit integer that specifies the type of data that is to be received from the secondary data source. Table 2-3 lists valid values for this field.

Table 2-3 Input types

Input identifier	Value	Description
kTrackModifierTypeMatrix	1	A 3x3 transformation matrix to transform the track's location, scaling, and so on.
kTrackModifierTypeClip	2	A QuickDraw clipping region to change the track's shape.
kTrackModifierTypeVolume	3	An 8.8 fixed-point value indicating the relative sound volume. This is used for fading the volume.
kTrackModifierTypeBalance	4	A 16-bit integer indicating the sound balance level. This is used for panning the sound location.
kTrackModifierTypeGraphicsMode	5	A graphics mode record (32-bit integer indicating graphics mode, followed by an RGB color) to modify the track's graphics mode for visual fades.
kTrackModifierObjectMatrix	6	A 3x3 transformation matrix to transform an object within the track's location, scaling, and so on.
kTrackModifierObjectGraphicsMode	7	A graphics mode record (32-bit integer indicating graphics mode, followed by an RGB color) to modify an object within the track's graphics mode for visual fades.
kTrackModifierTypeImage	'vide'	Compressed image data for an object within the track. Note that this was kTrackModifierTypeSpriteImage.

If the input is operating on an object within the track (for example, a sprite within a sprite track), an object ID atom must be included in the track input atom to identify the object.

The object ID atom contains the following data elements.

Field descriptions

Size	A 32-bit integer that specifies the number of bytes in this object ID atom.
Type	A 32-bit integer that identifies the atom type; this field must be set to 'obid'.
Object ID	A 32-bit integer identifying the object.


Media Atoms

Media atoms define a track’s movie data. The media atom contains information that specifies the media handler component that is to interpret the media data, and it also specifies the data references.

The media atom has an atom type of 'mdia'. It must contain a media header atom ('mdhd'), and it can contain a handler reference ('hdlr'), media information ('minf'), and user data ('udta').

Figure 2-15 shows the layout of a media atom.

Figure 2-15 The layout of a media atom



Required atom

Media atoms contain the following data elements.

Field descriptions

Size	A 32-bit integer that specifies the number of bytes in this media atom.
Type	A 32-bit integer that identifies the atom type; this field must be set to 'mdia'.

Media header atom

This atom contains the standard media information. See “Media Header Atoms” (page 68).

Handler reference atom

This atom identifies the media handler component that is to be used to interpret the media data. See “Handler Reference Atoms” (page 70) for more information.

Note that the handler reference atom tells you the kind of media this media atom contains—for example, video or sound. The layout of the media information atom is specific to the media handler that is to interpret the media. “Media Information Atoms” (page 72) discusses how data may be stored in a media, using the video media format defined by Apple as an example.

Media information atom

This atom contains data specific to the media type for use by the media handler component. See “Media Information Atoms” (page 72).

User data atom See “User Data Atoms” (page 43).

Media Header Atoms

The media header atom specifies the characteristics of a media, including time scale and duration. The media header atom has an atom type of 'mdhd'.

Figure 2-16 shows the layout of the media header atom.

Figure 2-16 The layout of a media header atom

Bytes	
Media header atom	
Atom size	4
Type = 'mdhd '	4
Version	1
Flags	3
Creation time	4
Modification time	4
Time scale	4
Duration	4
Language	2
Quality	2

Media header atoms contain the following data elements.

Field descriptions

Size	A 32-bit integer that specifies the number of bytes in this media header atom.
Type	A 32-bit integer that identifies the atom type; this field must be set to 'mdhd '.
Version	One byte that specifies the version of this header atom.
Flags	Three bytes of space for media header flags. Set this field to 0.
Creation time	A 32-bit integer that specifies (in seconds since midnight, January 1, 1904) when the media atom was created.
Modification time	A 32-bit integer that specifies (in seconds since midnight, January 1, 1904) when the media atom was changed.
Time scale	A time value that indicates the time scale for this media—that is, the number of time units that pass per second in its time coordinate system.

Movie Atoms

Duration	The duration of this media in units of its time scale.
Language	A 16-bit integer that specifies the language code for this media. See “Language Code Values” (page 253) for valid language codes.
Quality	A 16-bit integer that specifies the media’s playback quality—that is, its suitability for playback in a given environment.

Handler Reference Atoms

The handler reference atom specifies the media handler component that is to be used to interpret the media’s data. The handler reference atom has an atom type value of `'hdlr'`.

Historically, the handler reference atom was also used for data references. However, this use may now be ignored.

The handler atom within a media atom declares the process by which the media data in the stream may be presented, and thus, the nature of the media in a stream. For example, a video handler would handle a video track.

Figure 2-17 shows the layout of a handler reference atom.

Figure 2-17 The layout of a handler reference atom

Bytes	
Handler reference atom	
Atom size	4
Type = 'hdlr'	4
Version	1
Flags	3
Component type	4
Component subtype	4
Component manufacturer	4
Component flags	4
Component flags mask	4
Component name	Variable

Handler reference atoms contain the following data elements.

Field descriptions

Size	A 32-bit integer that specifies the number of bytes in this handler reference atom.
Type	A 32-bit integer that identifies the atom type; this field must be set to 'hdlr'.
Version	A 1-byte specification of the version of this handler information.
Flags	A 3-byte space for handler information flags. Set this field to 0.
Component type	A four-character code that identifies the type of the handler. Only two values are valid for this field: 'mhlr' for media handlers and 'dhlr' for data handlers.
Component subtype	

Movie Atoms

A four-character code that identifies the type of the media handler or data handler. For media handlers, this field defines the type of data—for example, 'vide' for video data or 'soun' for sound data.

For data handlers, this field defines the data reference type—for example, a component subtype value of 'alis' identifies a file alias.

Component manufacturer

Reserved. Set to 0.

Component flags

Reserved. Set to 0.

Component flags mask

Reserved. Set to 0.

Component name

A (counted) string that specifies the name of the component—that is, the media handler used when this media was created. This field may contain a zero-length (empty) string.

Media Information Atoms

Media information atoms (defined by the 'minf' atom type) store handler-specific information for a track's media data. The media handler uses this information to map from media time to media data and to process the media data.

These atoms contain information that is specific to the type of data defined by the media. Further, the format and content of media information atoms are dictated by the media handler that is responsible for interpreting the media data stream. Another media handler would not know how to interpret this information.

This section describes the atoms that store media information for the video ('vmhd'), sound ('smhd'), and base ('gmhd') portions of QuickTime movies.


Note

“Using Sample Atoms” (page 99) discusses how the video media handler locates samples in a video media. ♦

Video Media Information Atoms

Video media information atoms are the highest-level atoms in video media. These atoms contain a number of other atoms that define specific characteristics of the video media data. Figure 2-18 shows the layout of a video media information atom.

Figure 2-18 The layout of a media information atom for video



Required atom

The video media information atom contains the following data elements.

Field descriptions

Size A 32-bit integer that specifies the number of bytes in this video media information atom.

Type A 32-bit integer that identifies the atom type; this field must be set to 'minf'.

Video media information atom See “Video Media Information Header Atoms” (page 74).

Handler reference atom See “Handler Reference Atoms” (page 70).

Data information atom See “Data Information Atoms” (page 80).

Sample table atom See “Sample Table Atoms” (page 85).

Video Media Information Header Atoms

Video media information header atoms define specific color and graphics mode information.

Figure 2-19 shows the structure of a video media information header atom.

Figure 2-19 The layout of a media information header atom for video

Bytes	
Video media information header atom	
Atom size	4
Type = 'vmhd'	4
Version	1
Flags	3
Graphics mode	2
Opcolor	6

The video media information header atom contains the following data elements.

Field descriptions

Size	A 32-bit integer that specifies the number of bytes in this video media information header atom.
Type	A 32-bit integer that identifies the atom type; this field must be set to 'vmhd'.
Version	A 1-byte specification of the version of this video media information header atom.
Flags	A 3-byte space for video media information flags. There is one defined flag. No lean ahead

This is a compatibility flag that allows QuickTime to distinguish between movies created with QuickTime 1.0 and newer movies. You should always set this flag to 1, unless you are creating a movie intended for playback using version 1.0 of QuickTime. This flag's value is 0x0001.

- Graphics mode


A 16-bit integer that specifies the transfer mode. The transfer mode specifies which Boolean operation QuickDraw should perform when drawing or transferring an image from one location to another. See “Graphics Modes” (page 257) for a list of graphics modes supported by QuickTime.
- Opcolor

Three 16-bit values that specify the red, green, and blue colors for the transfer mode operation indicated in the graphics mode field.

Sound Media Information Atoms

Sound media information atoms are the highest-level atoms in sound media. These atoms contain a number of other atoms that define specific characteristics of the sound media data. Figure 2-20 shows the layout of a sound media information atom.

Figure 2-20 The layout of a media information atom for sound



Required atom

The sound media information atom contains the following data elements.

Field descriptions

Size A 32-bit integer that specifies the number of bytes in this sound media information atom.

Type A 32-bit integer that identifies the atom type; this field must be set to 'minf'.

Sound media information header atom See “Sound Media Information Header Atoms” (page 76).

Handler reference atom See “Handler Reference Atoms” (page 70).

Data information atom See “Data Information Atoms” (page 80).

Sample table atom See “Sample Table Atoms” (page 85).

Sound Media Information Header Atoms

The sound media information header atom (shown in Figure 2-21) stores the sound media’s control information, such as balance.

Figure 2-21 The layout of a sound media information header atom

Sound media information header atom		Bytes
Atom size		4
Type = 'smhd'		4
Version		1
Flags		3
Balance		2
Reserved		2

Movie Atoms

The sound media information header atom contains the following data elements.

Field descriptions


Size	A 32-bit integer that specifies the number of bytes in this sound media information header atom.
Type	A 32-bit integer that identifies the atom type; this field must be set to 'smhd'.
Version	A 1-byte specification of the version of this sound media information header atom.
Flags	A 3-byte space for sound media information flags. Set this field to 0.
Balance	A 16-bit integer that specifies the sound balance of this sound media. Sound balance is the setting that controls the mix of sound between the two speakers of a computer. This field is normally set to 0. See “Balance” (page 259) for more information about balance values.
Reserved	Reserved for use by Apple. Set this field to 0.

Base Media Information Atoms

The base media information atom (shown in Figure 2-22) stores the media information for media types such as text, MPEG, time code, and music.

Media types that are derived from the base media handler may add other atoms within the base media information atom, as appropriate. At present, the only media type that defines any additional atoms is the timecode media. See “Timecode Media Information Atom” (page 133) for more information about timecode media.

Figure 2-22 The layout of a base media information atom



Required atom

The base media information atom contains the following data elements.

Field descriptions

Size A 32-bit integer that specifies the number of bytes in this base media information atom.

Type A 32-bit integer that identifies the atom type; this field must be set to 'minf'.

Base media information header atom
See “Base Media Information Header Atoms” (page 78).

Base media info atom
See “Base Media Info Atoms” (page 79).

Base Media Information Header Atoms

The base media information header atom indicates that this media information atom pertains to a base media.

The base media information header atom contains the following data elements.

Field descriptions

Size A 32-bit integer that specifies the number of bytes in this base media information header atom.

Type A 32-bit integer that identifies the atom type; this field must be set to 'gmhd'.

Base Media Info Atoms

The base media info atom, contained in the base media information atom, defines the media’s control information, including graphics mode and balance information.

Figure 2-23 shows the layout of the base media info atom.

Figure 2-23 The layout of a base media info atom

Base media info atom	Bytes
Atom size	4
Type = 'gmin'	4
Version	1
Flags	3
Graphics mode	2
Opcolor	6
Balance	2
Reserved	2

The base media info atom contains the following data elements.

Field descriptions

Size	A 32-bit integer that specifies the number of bytes in this base media info atom.
Type	A 32-bit integer that identifies the atom type; this field must be set to 'gmin'.
Version	A 1-byte specification of the version of this base media information header atom.
Flags	A 3-byte space for base media information flags. Set this field to 0.
Graphics mode	A 16-bit integer that specifies the transfer mode. The transfer mode specifies which Boolean operation

Movie Atoms


	QuickDraw should perform when drawing or transferring an image from one location to another. See “Graphics Modes” (page 257) for more information about graphics modes supported by QuickTime.
Opcolor	Three 16-bit values that specify the red, green, and blue colors for the transfer mode operation indicated in the graphics mode field.
Balance	A 16-bit integer that specifies the sound balance of this media. Sound balance is the setting that controls the mix of sound between the two speakers of a computer. This field is normally set to 0. See “Balance” (page 259) for more information about balance values.
Reserved	Reserved for use by Apple. Set this field to 0.

Data Information Atoms

The handler reference atom (described in “Handler Reference Atoms” (page 70)) contains information specifying the data handler component that provides access to the media data. The data handler component uses the data information atom to interpret the media’s data. Data information atoms have an atom type value of ‘dinf’.

Figure 2-24 shows the layout of the data information atom.

Figure 2-24 The layout of a data information atom



The data information atom contains the following data elements.

Field descriptions

Size	A 32-bit integer that specifies the number of bytes in this data information atom.
Type	A 32-bit integer that identifies the atom type; this field must be set to 'dinf'.

Data reference atom

See “Data Reference Atoms” (page 82).

Data Reference Atoms

Data reference atoms contain tabular data that instructs the data handler component how to access the media’s data. Figure 2-24 shows the data reference atom.

The data reference atom contains the following data elements.

Field descriptions

Size	A 32-bit integer that specifies the number of bytes in this data reference atom.
Type	A 32-bit integer that identifies the atom type; this field must be set to 'dref'.
Version	A 1-byte specification of the version of this data reference atom.
Flags	A 3-byte space for data reference flags. Set this field to 0.
Number of entries	A 32-bit integer containing the count of data references that follow.
Data references	An array of data references.

Each data reference is formatted like an atom and contains the following data elements.

Size	A 32-bit integer that specifies the number of bytes in the data reference.
Type	A 32-bit integer that specifies the type of the data in the data reference. Table 2-4 lists valid type values.
Version	A 1-byte specification of the version of the data reference.

Movie Atoms

Flags	A 3-byte space for data reference flags. There is one defined flag. Self reference This flag indicates that the media’s data is in the same file as the movie atom. On the Macintosh, and other file systems with multifork files, set this flag to 1 even if the data resides in a different fork from the movie atom. This flag’s value is 0x0001.
Data	The data reference information.

Table 2-4 shows the currently defined data reference types that may be stored in a header atom.

Table 2-4 Data reference types


Data reference type	Description
'alis'	Data reference is a Macintosh alias. An alias contains information about the file, including its full path name.
'rsrc'	Data reference is a Macintosh alias. Appended to the end of the alias is the resource type (stored as a 32-bit integer) and ID (stored as a 16-bit signed integer) to use within the specified file.
'url '	A C string that specifies a URL. There may be additional data after the C string.

Sample Atoms

QuickTime stores media data in samples. A **sample** is a single element in a sequence of time-ordered data. Samples are stored in the media, and they may have varying durations.

Samples are stored in a series of **chunks** in a media. Chunks are a collection of data samples in a media that allow optimized data access. A chunk may contain one or more samples. Chunks in a media may have different sizes, and the individual samples within a chunk may have different sizes from one another, as shown in Figure 2-25.

Figure 2-25 Samples in a media



One way to describe a sample is to use a sample table atom. The sample table atom acts as a storehouse of information about the samples and contains a number of different types of atoms. The various atoms contain information that allows the media handler to parse the samples in the proper order. This approach enforces an ordering of the samples without requiring that the sample data be stored sequentially with respect to movie time in the actual data stream.

The next section discusses the sample table atom. Subsequent sections discuss each of the atoms that may reside in a sample table atom.

Sample Table Atoms

The sample table atom contains information for converting from media time to sample number to sample location. This atom also indicates how to interpret the sample (for example, whether to decompress the video data and, if so, how). This section describes the format and content of the sample table atom.

The sample table atom has an atom type of 'stbl'. It can contain the sample description atom, the time-to-sample atom, the sync sample atom, the sample-to-chunk atom, the sample size atom, the chunk offset atom, and the shadow sync atom.

The sample table atom contains all the time and data indexing of the media samples in a track. Using tables, it is possible to locate samples in time, determine their type, and determine their size, container, and offset into that container.


If the track that contains the sample table atom references no data, then the sample table atom does not need to contain any child atoms (not a very useful media track).

If the track that the sample table atom is contained in *does* reference data, then the following child atoms are required: sample description, sample size, sample to chunk, and chunk offset. All of the subtables of the sample table use the same total sample count.

The sample description atom must contain at least one entry. A sample description atom is required because it contains the data reference index field that indicates which data reference atom to use to retrieve the media samples. Without the sample description, it is not possible to determine where the media samples are stored. The sync sample atom is optional. If the sync sample atom is not present, all samples are implicitly sync samples.

Figure 2-26 shows the layout of the sample table atom.

Figure 2-26 The layout of a sample table atom



The sample table atom contains the following data elements.

Field descriptions


- Size** A 32-bit integer that specifies the number of bytes in this sample table atom.
- Type** A 32-bit integer that identifies the atom type; this field must be set to 'stbl'.
- Sample description atom** See “Sample Description Atoms” (page 87).
- Time-to-sample atom** See “Time-to-Sample Atoms” (page 88).
- Sync sample atom** See “Sync Sample Atoms” (page 91).
- Sample-to-chunk atom** See “Sample-to-Chunk Atoms” (page 93).
- Sample size atom** See “Sample Size Atoms” (page 96).
- Chunk offset atom** See “Chunk Offset Atoms” (page 97).
- Shadow sync atom** Reserved for future use.

Sample Description Atoms

The sample description atom stores information that allows you to decode samples in the media. The data stored in the sample description varies, depending on the media type. For example, in the case of video media, the sample descriptions are image description structures. The sample description information for each media type is explained in Chapter 3, “Media Data Atom Types.”

Figure 2-27 shows the layout of the sample description atom.

Figure 2-27 The layout of a sample description atom



The sample description atom has an atom type of 'stsd'. The sample description atom contains a table of sample descriptions. A media may have one or more sample descriptions, depending upon the number of different encoding schemes used in the media and on the number of files used to store the data. The sample-to-chunk atom identifies the sample description for each sample in the media by specifying the index into this table for the appropriate description (see “Sample-to-Chunk Atoms” (page 93)).

The sample description atom contains the following data elements.

Field descriptions

Size A 32-bit integer that specifies the number of bytes in this sample description atom.

Movie Atoms

Type	A 32-bit integer that identifies the atom type; this field must be set to 'std'.
Version	A 1-byte specification of the version of this sample description atom.
Flags	A 3-byte space for sample description flags. Set this field to 0.
Number of entries	A 32-bit integer containing the number of sample descriptions that follow.
Sample description table	An array of sample descriptions.

While the exact format of the sample description varies by media type, the first four fields of every sample description are the same. See Chapter 3, “Media Data Atom Types,” for details on various media.

Field descriptions

Sample description size	A 32-bit integer indicating the number of bytes in the sample description.
Data format	A 32-bit integer indicating the format of the stored data. This depends on the media type, but is usually either the compression format or the media type.
Reserved	Six bytes that must be set to 0.
Data reference index	A 16-bit integer that contains the index of the data reference to use to retrieve data associated with samples that use this sample description. Data references are stored in data reference atoms.

Time-to-Sample Atoms

Time-to-sample atoms store duration information for a media’s samples, providing a mapping from a time in a media to the corresponding data sample. The time-to-sample atom has an atom type of 'stts'.

You can determine the appropriate sample for any time in a media by examining the time-to-sample atom table, which is contained in the time-to-sample atom.


Movie Atoms

The atom contains a compact version of a table that allows indexing from time to sample number. Other tables provide sample sizes and pointers from the sample number. Each entry in the table gives the number of consecutive samples with the same time delta, and the delta of those samples. By adding the deltas, a complete time-to-sample map can be built.

The atom contains time deltas: $DT(n+1) = DT(n) + STTS(n)$ where $STTS(n)$ is the (uncompressed) table entry for sample n and DT is the display time for sample (n) . The sample entries are ordered by time stamps; therefore, the deltas are all nonnegative. The DT axis has a zero origin; $DT(i) = \text{SUM (for } j=0 \text{ to } i-1 \text{ of } \text{delta}(j))$, and the sum of all deltas gives the length of the media in the track (not mapped to the overall time scale, and not considering any edit list). The edit list atom provides the initial DT value if it is nonempty (nonzero).

Figure 2-28 shows the layout of the time-to-sample atom.

Figure 2-28 The layout of a time-to-sample atom



The time-to-sample atom contains the following data elements.

Field descriptions

- Size** A 32-bit integer that specifies the number of bytes in this time-to-sample atom.
- Type** A 32-bit integer that identifies the atom type; this field must be set to 'stts'.

Version	A 1-byte specification of the version of this time-to-sample atom.
Flags	A 3-byte space for time-to-sample flags. Set this field to 0.
Number of entries	A 32-bit integer containing the count of entries in the time-to-sample table.
Time-to-sample table	A table that defines the duration of each sample in the media. Each table entry contains a count field and a duration field. The structure of the time-to-sample table is shown in Figure 2-29.

Figure 2-29 The layout of a time-to-sample table entry

Sample count	Sample duration	Field
4	4	Bytes

You define a time-to-sample table entry by specifying these fields:

Field descriptions

Sample count	A 32-bit integer that specifies the number of consecutive samples that have the same duration.
Sample duration	A 32-bit integer that specifies the duration of each sample.

Entries in the table describe samples according to their order in the media and their duration. If consecutive samples have the same duration, a single table entry can be used to define more than one sample. In these cases, the count field indicates the number of consecutive samples that have the same duration. For example, if a video media has a constant frame rate, this table would have one entry and the count would be equal to the number of samples.

Figure 2-30 presents an example of a time-to-sample table that is based on the chunked media data shown in Figure 2-25 (page 84). That data stream contains a total of nine samples that correspond in count and duration to the entries of the table shown here. Even though samples 4, 5, and 6 are in the same chunk, sample 4 has a duration of 3, and samples 5 and 6 have a duration of 2.

Figure 2-30 An example of a time-to-sample table

count	duration
4	3
2	1
3	2

Sync Sample Atoms


The sync sample atom identifies the **key frames** in the media. In a media that contains compressed data, key frames define starting points for portions of a temporally compressed sequence. The key frame is self-contained—that is, it is independent of preceding frames. Subsequent frames may depend on the key frame.

The sync sample atom provides a compact marking of the random access points within a stream. The table is arranged in strictly increasing order of sample number. If this table is not present, every sample is implicitly a random access point.

Sync sample atoms have an atom type of 'stss'. The sync sample atom contains a table of sample numbers. Each entry in the table identifies a sample that is a key frame for the media. If no sync sample atom exists, then all the samples are key frames.

Figure 2-31 shows the layout of a sync sample atom.

Figure 2-31 The layout of a sync sample atom



The sync sample atom contains the following data elements.

Field descriptions

Size	A 32-bit integer that specifies the number of bytes in this sync sample atom.
Type	A 32-bit integer that identifies the atom type; this field must be set to 'stss'.
Version	A 1-byte specification of the version of this sync sample atom.
Flags	A 3-byte space for sync sample flags. Set this field to 0.
Number of entries	A 32-bit integer containing the count of entries in the sync sample table.
Sync sample table	A table of sample numbers; each sample number corresponds to a key frame. Figure 2-32 shows the layout of the sync sample table.

Figure 2-32 The layout of a sync sample table

Number	Sample 1
Number	Sample 2
Number	Sample 3
Number	Sample 4
Number	Sample 5


Sample-to-Chunk Atoms

As samples are added to a media, they are collected into chunks that allow optimized data access. A chunk contains one or more samples. Chunks in a media may have different sizes, and the samples within a chunk may have different sizes. The sample-to-chunk atom stores chunk information for the samples in a media.

Sample-to-chunk atoms have an atom type of `'stsc'`. The sample-to-chunk atom contains a table that maps samples to chunks in the media data stream. By examining the sample-to-chunk atom, you can determine the chunk that contains a specific sample.

Figure 2-33 shows the layout of the sample-to-chunk atom.

Figure 2-33 The layout of a sample-to-chunk atom



The sample-to-chunk atom contains the following data elements.

Field descriptions

- Size** A 32-bit integer that specifies the number of bytes in this sample-to-chunk atom.
- Type** A 32-bit integer that identifies the atom type; this field must be set to 'stsc'.
- Version** A 1-byte specification of the version of this sample-to-chunk atom.
- Flags** A 3-byte space for sample-to-chunk flags. Set this field to 0.
- Number of entries** A 32-bit integer containing the count of entries in the sample-to-chunk table.

Sample-to-chunk table

A table that maps samples to chunks. Figure 2-34 shows the structure of an entry in a sample-to-chunk table. Each sample-to-chunk atom contains such a table, which identifies the chunk for each sample in a media. Each entry in the table contains a first chunk field, a samples per chunk field, and a sample description ID field. From this information, you can ascertain where samples reside in the media data.

Figure 2-34 The layout of a sample-to-chunk table entry

First chunk	Samples per chunk	Sample description ID	Fields
4	4	4	Bytes

You define a sample-to-chunk table entry by specifying the following data elements.

Field descriptions

First chunk The first chunk number using this table entry.

Samples per chunk The number of samples in each chunk.

Sample description ID

The identification number associated with the sample description for the sample. For details on sample description atoms, see “Sample Description Atoms” (page 87).

Figure 2-35 shows an example of a sample-to-chunk table that is based on the data stream shown in Figure 2-25 (page 84).

Figure 2-35 An example of a sample-to-chunk table

First chunk	Samples per chunk	Sample description ID
1	3	23
3	1	23
5	1	24

Each table entry corresponds to a set of consecutive chunks, each of which contains the same number of samples. Furthermore, each of the samples in

these chunks must use the same sample description. Whenever the number of samples per chunk or the sample description changes, you must create a new table entry. If all the chunks have the same number of samples per chunk and use the same sample description, this table has one entry.


Sample Size Atoms

You use sample size atoms to specify the size of each sample in the media. Sample size atoms have an atom type of 'stsz'.

The sample size atom contains the sample count and a table giving the size of each sample. This allows the media data itself to be unframed. The total number of samples in the media is always indicated in the sample count. If the default size is indicated, then no table follows.

Figure 2-36 shows the layout of the sample size atom.

Figure 2-36 The layout of a sample size atom



The sample size atom contains the following data elements.

Field descriptions

Size A 32-bit integer that specifies the number of bytes in this sample size atom.

Movie Atoms

Type	A 32-bit integer that identifies the atom type; this field must be set to 'stsz'.
Version	A 1-byte specification of the version of this sample size atom.
Flags	A 3-byte space for sample size flags. Set this field to 0.
Sample size	A 32-bit integer specifying the sample size. If all the samples are the same size, this field contains that size value. If this field is set to 0, then the samples have different sizes, and those sizes are stored in the sample size table.
Number of entries	A 32-bit integer containing the count of entries in the sample size table.
Sample size table	A table containing the sample size information. The sample size table contains an entry for every sample in the media's data stream. Each table entry contains a size field. The size field contains the size, in bytes, of the sample in question. The table is indexed by sample number—the first entry corresponds to the first sample, the second entry is for the second sample, and so on.

Figure 2-37 shows a sample size table.

Figure 2-37 An example of a sample size table

Size	Sample 1
Size	Sample 2
Size	Sample 3
Size	Sample 4
Size	Sample 5

Chunk Offset Atoms

Chunk offset atoms identify the location of each chunk of data in the media's data stream. Chunk offset atoms have an atom type of 'stco'.

The chunk-offset table gives the index of each chunk into the containing file. There are two variants, permitting the use of 32-bit or 64-bit offsets. The latter is useful when managing very large movies. Only one of these variants occurs in any single instance of a sample table atom.


Note that offsets are file offsets, not the offset into any atom within the file (for example, a 'mdat' atom). This permits referring to media data in files without any atom structure. However, be careful when constructing a self-contained QuickTime file with its meta-data (movie atom) at the front because the size of the movie atom affects the chunk offsets to the media data.

Note

The sample table atom can contain a 64-bit chunk offset atom (STChunkOffset64AID = 'co64'). When this atom appears, it is used in place of the original chunk offset atom, which can contain only 32-bit offsets. When QuickTime writes movie files, it uses the 64-bit chunk offset atom only if there are chunks that use the high 32-bits of the chunk offset. Otherwise, the original 32-bit chunk offset atom is used to ensure compatibility with previous versions of QuickTime. ♦

Figure 2-38 shows the layout of a chunk offset atom.

Figure 2-38 The layout of a chunk offset atom



The chunk offset atom contains the following data elements.

Field descriptions

Size	A 32-bit integer that specifies the number of bytes in this chunk offset atom.
Type	A 32-bit integer that identifies the atom type; this field must be set to 'stco'.
Version	A 1-byte specification of the version of this chunk offset atom.
Flags	A 3-byte space for chunk offset flags. Set this field to 0.
Number of entries	A 32-bit integer containing the count of entries in the chunk offset table.
Chunk offset table	A chunk offset table consisting of an array of offset values. There is one table entry for each chunk in the media. The offset contains the byte offset from the beginning of the data stream to the chunk. The table is indexed by chunk number—the first table entry corresponds to the first chunk, the second table entry is for the second chunk, and so on.

Figure 2-39 shows an example of a chunk offset table.

Figure 2-39 An example of a chunk offset table

Offset	Chunk 1
Offset	Chunk 2
Offset	Chunk 3
Offset	Chunk 4
Offset	Chunk 5

Using Sample Atoms

This section presents examples using the atoms just described. These examples are intended to help you understand the relationships between these atoms.

The first section, “Finding a Sample” (page 100), describes the steps that the video media handler uses to find the sample that contains the media data for a particular time in a media. The second section, “Finding a Key Frame” (page 100), describes the steps that the video media handler uses to find an appropriate key frame for a specific time in a movie.

Finding a Sample

When QuickTime displays a movie or track, it “tells” the appropriate media handler to access the media data for a particular time. The media handler must correctly interpret the data stream to retrieve the requested data. In the case of video media, the media handler traverses several atoms to find the location and size of a sample for a given media time.

The media handler performs the following steps:

1. Determines the time in the media time coordinate system.
2. Examines the time-to-sample atom to determine the sample number that contains the data for the specified time.
3. Scans the sample-to-chunk atom to discover which chunk contains the sample in question.
4. Extracts the offset to the chunk from the chunk offset atom.
5. Finds the offset within the chunk and the sample’s size by using the sample size atom.

Finding a Key Frame

Finding a key frame for a specified time in a movie is slightly more complicated than finding a sample for a specified time. The media handler must use the sync sample atom and the time-to-sample atom together in order to find a key frame.

The media handler performs the following steps:

1. Examines the time-to-sample atom to determine the sample number that contains the data for the specified time.
2. Scans the sync sample atom to find the key frame that precedes the sample number chosen in step 1.
3. Scans the sample-to-chunk atom to discover which chunk contains the key frame.

4. Extracts the offset to the chunk from the chunk offset atom.
5. Finds the offset within the chunk and the sample's size by using the sample size atom.

Compressed Movie Resources

Most QuickTime movies have meta-data in addition to their media data. **Media data** can be compressed using a variety of video and sound compression algorithms. Beginning with QuickTime 3, it also became possible to compress the **meta-data**—more commonly known as the **movie resource**. However, the movie resource cannot be compressed by means of a lossy compression algorithm because it contains critical information, such as the video and audio compression types used, individual frame offsets, and timing information. To compress the movie resource, therefore, *lossless* data compression algorithms must be used.

Compressing movie resources using data compression typically reduces the size of the movie resource by 50% or more. For QuickTime movies that are streamed over the Internet, this can substantially reduce the startup latency of the movie, and therefore has a number of distinct advantages.

Allowing QuickTime to Compress the Movie Resource

Most application developers won't need to know the details of how movie resources are compressed. The Movie Toolbox `FlattenMovie` and `FlattenMovieData` functions compress the movie resource if so requested by the application. To accomplish this, applications only need to set the `flattenCompressMovieResource` flag when calling either function. The QuickTime movie export component also provides users with the option of compressing the movie resource when exporting or creating a new movie through export.

Structure of a Compressed Movie Resource

A compressed movie resource, similar to an uncompressed movie resource, is made up of a group of QuickTime atoms arranged in a hierarchy.

Like an uncompressed movie resource, the outermost atom is a movie atom. Within the movie atom, there is a single compressed movie atom, which

contains all other required atoms. The compressed movie atom has two subatoms. The first is a data compression atom, which contains a single 32-bit integer that identifies what lossless data compression algorithm was used to compress the movie resource. The second child atom is the compressed movie data, which contains the compressed movie resource itself. The first 32-bit integer in the compressed movie data atom indicates the uncompressed size of the movie resource, and then the compressed movie resource data follows.

The contents of a complete compressed movie are shown in Table 2-5. The constants that define the atom types are defined in `MoviesFormat.h`. The four-character codes for each atom type are also shown.

Table 2-5 Contents of complete compressed movie

Atom type	Four-character code
Movie	'moov'
Compressed movie	'cmov'
Data compression atom	'dcom'
Compressed movie data	'cmvd'
32-bit integer	Uncompressed size


Reference Movies

A QuickTime movie can act as a container for a set of alternate movies that should be displayed under specified conditions. One of these movies may be contained within the same file; any others are included by reference.

For example, a QuickTime movie can contain a list of references to movies having different data rates, allowing an application to choose the best-looking movie that can play smoothly as it downloads over the Internet, based on the user's connection speed.

A movie that contains references to alternate movies is called a reference movie.

A reference movie contains a reference movie atom ('rmra') at the top level of the movie atom. The movie atom may also contain a movie header atom, or it may contain the reference movie atom alone.

Figure 2-40 A movie atom containing a 'rmra' atom instead of a 'mvhd' atom

The reference movie atom contains one or more reference movie descriptor atoms, each of which describes an alternate movie.

Each reference movie descriptor atom contains a data reference atom, which specifies the location of a movie.

Note

Movie locations are specified using QuickTime data references. QuickTime supports multiple types of data reference, but alternate movies are generally specified using data reference types of either url ('url ') or file alias ('alis').

A reference movie descriptor atom may contain other atoms that specify the movie's system requirements and the movie quality. If so, there will be an atom of an appropriate type for each requirement that must be met for the movie to play, and there may be a quality atom as well.

Applications should play the highest-quality movie whose requirements are met by the user's system. If the data reference to the selected movie cannot be resolved—because the file cannot be found, for example—the application should recursively attempt to play the next-highest-quality movie until it succeeds or has exhausted the list of movies whose requirements are met.

If a movie contains both a reference movie atom and a movie header atom, applications should play the appropriate movie indicated by the reference movie atom.

If the user's system does not meet any of the alternate movies' criteria, or none of the qualifying data references can be resolved, applications should play the movie defined in the movie header atom. (The movie defined in the movie header atom can also be indicated by one of the alternate movie references.)

The movie header atom is sometimes used to provide a fallback movie for applications that can play older QuickTime movies but do not understand reference movies.

Reference Movie Atom

A reference movie atom contains references to one or more movies. It can optionally contain a list of system requirements in order for each movie to play, and a quality rating for each movie. It is typically used to specify a list of alternate movies to be played under different conditions.

A reference movie atom's parent is always a movie atom ('moov'). Only one reference movie atom is allowed in a given movie atom.

A reference movie atom may contain the following information.


Field descriptions

Size The number of bytes in this reference movie atom.

Type The type of this atom; this field must be set to 'rmra'.

Reference movie descriptor atom

A reference movie atom must contain at least one reference movie descriptor atom, and typically contains more than one. See "Reference Movie Descriptor Atom" (page 105) for more information.

Figure 2-41 A 'rmra' atom with multiple 'rmda' atoms

Reference Movie Descriptor Atom

Each reference movie descriptor atom contains other atoms that describe where a particular movie can be found, and optionally what the system requirements are to play that movie, as well as an optional quality rating for that movie.

A reference movie descriptor atom's parent is always a movie reference atom ('rmra'). Multiple reference movie descriptor atoms are allowed in a given movie reference atom, and more than one is usually present.

A reference movie descriptor atom may contain the following information.

Field descriptions


Size	The number of bytes in this reference movie descriptor atom.
Type	The type of this atom; this field must be set to 'rmda'.
Data reference atom	Each reference movie atom must contain exactly one data reference atom. See “Data Reference Atom” (page 106) for more information.
Data rate atom	A reference movie atom may contain an optional data rate atom. Only one data rate atom can be present. See “Data Rate Atom” (page 107) for more information.
CPU speed atom	A reference movie atom may contain an optional CPU speed atom. Only one CPU speed atom can be present. See “CPU Speed Atom” (page 108) for more information.

Version check atom A reference movie atom may contain an optional version check atom. Multiple version check atoms can be present. See “Version Check Atom” (page 109) for more information.

Component detect atom A reference movie atom may contain an optional component detect atom. Multiple component detect atoms can be present. See “Component Detect Atom” (page 110) for more information.

Quality atom A reference movie atom may contain an optional quality atom. Only one quality atom can be present. See “Quality Atom” (page 111) for more information.

Figure 2-42 Reference movie descriptor atom



Data Reference Atom

A data reference atom contains the information necessary to locate a movie, or a stream or file that QuickTime can play, typically in the form of a URL or a file alias.

Only one data reference atom is allowed in a given movie reference descriptor atom.

A data reference atom may contain the following information.

Field descriptions

Size	The number of bytes in this data reference atom.
Type	The type of this atom; this field must be set to 'rdrf'.
Flags	A 32-bit integer containing flags. One flag is currently defined: movie is self-contained.
	If the least-significant bit is set to 1, the movie is self-contained. This requires that the parent movie contain a movie header atom as well as a reference movie atom. In other words, the current 'moov' atom must contain both a 'rmra' atom and a 'mvhd' atom. To resolve this data reference, an application uses the movie defined in the movie header atom, ignoring the remainder of the fields in this data reference atom, which are used only to specify external movies.
Data reference type	The data reference type. A value of 'alis' indicates a file system alias record. A value of 'url' indicates a string containing a uniform resource locator. Note that the fourth character in 'url' is an ASCII blank (hex 20).
Data reference size	The size of the data reference in bytes, expressed as a 32-bit integer.
Data reference	A data reference to a QuickTime movie, or to a stream or file that QuickTime can play. If the reference type is 'alis' this field contains the contents of an <code>AliasHandle</code> . If the reference type is 'url' this field contains a null-terminated string that can be interpreted as a URL. The URL can be absolute or relative, and can specify any protocol that QuickTime supports, including <code>http://</code> , <code>ftp://</code> , <code>rtsp://</code> , <code>file:///</code> , and <code>data:</code> .

Data Rate Atom

A data rate atom specifies the minimum data rate required to play a movie. This is normally compared to the connection speed setting in the user's QuickTime Settings control panel. Applications should play the movie with the highest data rate less than or equal to the user's connection speed. If the connection speed is slower than any movie's data rate, applications should play the movie

with the lowest data rate. The movie with the highest data rate is assumed to have the highest quality.

Only one data rate atom is allowed in a given reference movie descriptor atom.

A data rate atom may contain the following information.

Field descriptions

Size	The number of bytes in this data rate atom.
Type	The type of this atom; this field must be set to 'rmdr'.
Flags	A 32-bit integer that is currently always 0.
Data rate	The required data rate in bits per second, expressed as a 32-bit integer.

CPU Speed Atom

A CPU speed atom specifies the minimum computing power needed to display a movie. QuickTime performs an internal test to determine the speed of the user's computer.

This is not a simple measurement of clock speed—it is a measurement of performance for QuickTime-related operations. Speed is expressed as a relative value between 100 and 2^{31} , in multiples of 100.

Note

Typical scores might range from a minimum score of 100, which would describe a computer as slow as, or slower than, a 166 MHz Pentium or 120 MHz PowerPC, to a maximum score of 600 for a 500 MHz Pentium III or 400 MHz G4 PowerPC. A computer with a graphics accelerator and a Gigahertz clock speed might score as high as 1000. Future computers will score higher.

Applications should play the movie with the highest specified CPU speed that is less than or equal to the user's speed. If the user's speed is lower than any movie's CPU speed, applications should play the movie with the lowest CPU speed requirement. The movie with the highest CPU speed is assumed to be the highest quality.

Only one CPU speed atom is allowed in a given reference movie descriptor atom.

A CPU speed atom may contain the following information.

Movie Atoms

Field descriptions

Size	The number of bytes in this CPU speed atom.
Type	The type of this atom; this field must be set to 'rmcs'.
Flags	A 32-bit integer that is currently always 0.
CPU speed	A relative ranking of required computer speed, expressed as a 32-bit integer divisible by 100, with larger numbers indicating higher speed.

Version Check Atom

A version check atom specifies a software package, such as QuickTime or QuickTime VR, and the version of that package needed to display a movie. The package is specified using a Macintosh Gestalt type, such as 'qtim' for QuickTime (QuickTime provides support for these Gestalt tests in the Windows computing environment).

You can specify a minimum required version to be returned by the Gestalt check, or you can require that a specific value be returned after performing a binary AND operation on the Gestalt bitfield and a mask.

Multiple version check atoms are allowed within a given reference movie descriptor atom. Applications should not attempt to play a movie unless all version checks are successful.

A version check atom may contain the following information.

Field descriptions

Size	The number of bytes in this version check atom.
Type	The type of this atom; this field must be set to 'rmvc'.
Flags	A 32-bit integer that is currently always 0.
Software package	A 32-bit Gestalt type, such as 'qtim', specifying the software package to check for.
Version	An unsigned 32-bit integer containing either the minimum required version or the required value after a binary AND operation.
Mask	The mask for a binary AND operation on the Gestalt bitfield.
Check type	The type of check to perform, expressed as 16-bit integer. Set to 0 for a minimum version check, set to 1 for a required value after a binary AND of the Gestalt bitfield and the mask.

Component Detect Atom

A component detect atom specifies a QuickTime component, such as a particular video decompressor, required to play the movie. The component type, subtype, and other required attributes can be specified, as well as a minimum version.

Multiple component detect atoms are allowed within a given reference movie descriptor atom. Applications should not attempt to play a movie unless at least the minimum versions of all required components are present.

A component detect atom may contain the following information.

Field descriptions

Size	The number of bytes in this component detect atom.
Type	The type of this atom; this field must be set to 'rmcd'.
Flags	A 32-bit integer that is currently always 0.
Component description	A component description record. For details, see “Component Description Record” (page 110).
Minimum version	An unsigned 32-bit integer containing the minimum required version of the specified component.

Component Description Record

Describes a class of components by their attributes. Fields that are set to 0 are treated as “don’t care.”

```
struct ComponentDescription {
    OSType      componentType;
    OSType      componentSubType;
    OSType      componentManufacturer;
    unsigned long componentFlags;
    unsigned long componentFlagsMask;
};
```

<code>componentType</code>	A four-character code that identifies the type of component.
<code>componentSubType</code>	A four-character code that identifies the subtype of the component. For example, the subtype of an image compressor component indicates the compression

algorithm employed by the compressor. A value of 0 matches any subtype.

componentManufacturer

A four-character code that identifies the manufacturer of the component. Components provided by Apple have a manufacturer value of 'appl'. A value of 0 matches any manufacturer

componentFlags

A 32-bit field that contains flags describing required component capabilities. The high-order 8 bits should be set to 0. The low-order 24 bits are specific to each component type. These flags can be used to indicate the presence of features or capabilities in a given component.

componentFlagsMask

A 32-bit field that indicates which flags in the componentFlags field are relevant to this operation. For each flag in the componentFlags field that is to be considered as a search criterion, set the corresponding bit in this field to 1. To ignore a flag, set the bit to 0.

Constants

canMovieImportInPlace

Set this bit if a movie import component must be able to create a movie from a file without having to write to a separate disk file. Examples include MPEG and AIFF import components.

movieImportSubTypeIsFileExtension

Set this bit if the component's subtype is a file extension instead of a Macintosh file type. For example, if you require an import component that opens files with an extension of .doc, set this flag and set your component subtype to 'DOC'.

canMovieImportFiles

Set this bit if a movie import component must import files.

Quality Atom

A quality atom describes the relative quality of a movie. This acts as a tie-breaker if more than one movie meets the specified requirements, and it is not otherwise obvious which movie should be played.

Movie Atoms

This would be the case if two qualified movies have the same data rate and CPU speed requirements, for example, or if one movie requires a higher data rate and another requires a higher CPU speed, but both can be played on the current system. In these cases, applications should play the movie with the highest quality, as specified in the quality atom.

Only one quality atom is allowed in a given reference movie descriptor atom.

A quality atom may contain the following information.

Field descriptions

Size	The number of bytes in this quality atom.
Type	The type of this atom; this field must be set to 'rmqu'.
Quality	The relative quality of the movie, expressed as a 32-bit integer. A larger number indicates higher quality. A unique value should be given to each movie.

Media Data Atom Types

QuickTime uses atoms of different types to store different types of media data—video media atoms for video data, sound media atoms for audio data, and so on. This chapter discusses in detail each of these different media data atom types.

If you are a QuickTime application or tool developer, you'll want to read this chapter in order to understand the fundamentals of how QuickTime uses atoms for storage of different media data. The information in this chapter is current as of QuickTime 4.1. For the latest updates and postings on Apple's QuickTime developer website, be sure to see <http://developer.apple.com/techpubs/quicktime/quicktime.html>.

This chapter is divided into the following major sections:

- “Video Media” (page 114) describes video media, which is used to store compressed and uncompressed image data in QuickTime movies.
- “Sound Media” (page 124) discusses sound media used to store compressed and uncompressed audio data in QuickTime movies.
- “Timecode Media” (page 131) describes time code media used to store time code data in QuickTime movies.
- “Text Media” (page 134) discusses text media used to store text data in QuickTime movies.
- “Music Media” (page 139) discusses music media used to store note-based audio data, such as MIDI data, in QuickTime movies.
- “MPEG Media” (page 140) discusses MPEG media used to store MPEG streams in QuickTime movies.
- “Sprite Media” (page 140) discusses sprite media used to store character-based animation data in QuickTime movies.

- “Tween Media” (page 168) discusses tween media used to store pairs of values to be interpolated between in QuickTime movies.
- “Modifier Tracks” (page 180) discusses the capabilities of modifier tracks.
- “Track References” (page 181) describes a feature of QuickTime that allows you to relate a movie’s tracks to one another.
- “3D Media” (page 183) discusses briefly how QuickTime movies store 3D image data in a base media.
- “Hint Media” (page 185) describes the additions to the QuickTime file format for streaming QuickTime movies over the Internet.
- “VR Media” (page 203) describes the QuickTime VR world and node information atom containers, as well as cubic panoramas, which are new to QuickTime VR 3.0.
- “Movie Media” (page 246) discusses movie media which is used to encapsulate embedded movies within QuickTime movies.

Video Media

Video media is used to store compressed and uncompressed image data in QuickTime movies. It has a media type of 'vide'.

Video Sample Description

The video sample description contains information that defines how to interpret video media data. This sample description is based on the standard sample description, as described in “Sample Description Atoms” (page 87).

Media Data Atom Types

The data format field of a video sample description indicates the type of compression that was used to compress the image data. Table 3-1 shows some of the formats supported.

Table 3-1 Some image compression formats

Compression type	Description
'cvid'	Cinepak
'jpeg'	JPEG
'raw '	Uncompressed RGB
'Yuv2'	Uncompressed YUV422
'smc '	Graphics
'rle '	Animation
'rpza'	Apple video
'kpcd'	Kodak Photo CD
'mpeg'	MPEG
'mjpa'	Motion-JPEG (format A)
'mjpb'	Motion-JPEG (format B)
'svqi'	Sorenson video

The video media handler also adds some of its own fields to the sample description.

Field descriptions

Version	A 16-bit integer indicating the version number of the compressed data. This is set to 0, unless a compressor has changed its data format.
Revision level	A 16-bit integer that must be set to 0.
Vendor	A 32-bit integer that specifies the developer of the compressor that generated the compressed data. Often this field contains 'appl' to indicate Apple Computer, Inc.

Media Data Atom Types

Temporal quality	A 32-bit integer containing a value from 0 to 1023 indicating the degree of temporal compression.
Spatial quality	A 32-bit integer containing a value from 0 to 1024 indicating the degree of spatial compression.
Width	A 16-bit integer that specifies the width of the source image in pixels.
Height	A 16-bit integer that specifies the height of the source image in pixels.
Horizontal resolution	A 32-bit fixed-point number containing the horizontal resolution of the image in pixels per inch.
Vertical resolution	A 32-bit fixed-point number containing the vertical resolution of the image in pixels per inch.
Data size	A 32-bit integer that must be set to 0.
Frame count	A 16-bit integer that indicates how many frames of compressed data are stored in each sample. Usually set to 1.
Compressor name	A 32-byte Pascal string containing the name of the compressor that created the image, such as "jpeg".
Depth	A 16-bit integer that indicates the pixel depth of the compressed image. Values of 1, 2, 4, 8, 16, 24, and 32 indicate the depth of color images. The value 32 should be used only if the image contains an alpha channel. Values of 34, 36, and 40 indicate 2-, 4-, and 8-bit grayscale, respectively, for grayscale images.
Color table ID	A 16-bit integer that identifies which color table to use. If this field is set to -1, the default color table should be used for the specified depth. For all depths below 16 bits per pixel, this indicates a standard Macintosh color table for the specified depth. Depths of 16, 24, and 32 have no color table. If the color table ID is set to 0, a color table is contained within the sample description itself. The color table immediately follows the color table ID field in the sample description. See "Color Table Atoms" (page 42) for a complete description of a color table.

Video sample descriptions can be extended by appending other atoms. These atoms are placed after the color table, if one is present. These extensions to the

sample description may contain display hints for the decompressor or may simply carry additional information associated with the images. Table 3-2 lists the currently defined extensions to video sample descriptions.

Table 3-2 Video sample description extensions

Extension type	Description
'gama'	A 32-bit fixed-point number indicating the gamma level at which the image was captured. The decompressor can use this value to gamma-correct at display time.
'fiel'	Two 8-bit integers that define field handling. This information is used by applications to modify decompressed image data or by decompressor components to determine field display order. The first byte specifies the field count, and may be set to 1 or 2. When the field count is 2, the second byte specifies the field ordering: which contains the topmost scan-line, which field should be displayed earliest, and which is stored first in each sample. Each QuickTime sample consists of two distinct compressed images, each coding one field: the field with the topmost scan-line, T, and the other field, B. Each field is half the height of the overall image, as declared in the height field of the sample description. To be precise, if the height field contains the value H, then the field T has $((H+1) \text{ div } 2)$ lines, and field B has $(H \text{ div } 2)$ lines. The following defines the permitted variants: 0 – field ordering is unknown. 1 – T is displayed earliest, T is stored first in the file. 6 – B is displayed earliest, B is stored first in the file.
'mjqt'	The default quantization table for a Motion-JPEG data stream.
'mjht'	The default Huffman table for a Motion-JPEG data stream.

Video Sample Data

The format of the data stored in video samples is completely dependent on the type of the compressed data stored in the video sample description. The

following sections discuss each of the video encoding schemes supported by QuickTime.

Uncompressed RGB

Uncompressed RGB data is stored in a variety of different formats. The format used depends on the depth field of the video sample description. For all depths, the image data is padded on each scan line to ensure that each scan line begins on an even byte boundary.

- For depths of 1, 2, 4, and 8, the values stored are indexes into the color table specified in the color table ID field.
- For a depth of 16, the pixels are stored as 5-5-5 RGB values with the high bit of each 16-bit integer set to 0.
- For a depth of 24, the pixels are stored packed together in RGB order.
- For a depth of 32, the pixels are stored with an 8-bit alpha channel, followed by 8-bit RGB components.

Uncompressed yuv2

The yuv2 stream is encoded in a series of 4-byte packets. Each packet represents two adjacent pixels on the same scan line. The bytes within each packet are ordered as follows:

y0 u y1 v

y0 is the luminance value for the left pixel; y1 the luminance for the right pixel. u and v are chromatic values that are shared by both pixels. The conversion into RGB space is represented by the following equations:

$$r = 1.402 * v + y + .5$$

$$g = y - .7143 * v - .3437 * u + .5$$

$$b = 1.77 * u + y + .5$$

The *r*, *g*, and *b* values range from 0 to 255.

JPEG

QuickTime stores JPEG images according to the rules described in the ISO JPEG specification, document number DIS 10918-1.

Motion-JPEG

Motion-JPEG (M-JPEG) is a variant of the ISO JPEG specification for use with digital video streams. Instead of compressing an entire image into a single bitstream, Motion-JPEG compresses each video field separately, returning the resulting JPEG bitstreams consecutively in a single frame.

There are two flavors of Motion-JPEG currently in use. These two formats differ based on their use of markers. Motion-JPEG format A supports markers; Motion-JPEG format B does not. The following paragraphs describe how QuickTime stores Motion-JPEG sample data. Figure 3-1 shows an example of Motion-JPEG A dual-field sample data. Figure 3-2 shows an example of Motion-JPEG B dual-field sample data.

Each field of Motion-JPEG format A fully complies with the ISO JPEG specification, and therefore supports application markers. QuickTime uses the APP₁ marker to store control information, as follows (all of the fields are 32-bit integers):

Field descriptions

Reserved	Unpredictable; should be set to 0.
Tag	Identifies the data type; this field must be set to 'mjpg'.
Field size	The actual size of the image data for this field, in bytes.
Padded field size	Contains the size of the image data, including pad bytes. Some video hardware may append pad bytes to the image data; this field, along with the field size field, allows you to compute how many pad bytes were added.
Offset to next field	The offset, in bytes, from the start of the field data to the start of the next field in the bitstream. This field should be set to 0 in the last field's marker data.
Quantization table offset	The offset, in bytes, from the start of the field data to the quantization table marker. If this field is set to 0, check the image description for a default quantization table.
Huffman table offset	The offset, in bytes, from the start of the field data to the Huffman table marker. If this field is set to 0, check the image description for a default Huffman table.
Start Of Frame offset	

Media Data Atom Types

The offset from the start of the field data to the start of image marker. This field should never be set to 0.

Start Of Scan offset

The offset, in bytes, from the start of the field data to the start of the scan marker. This field should never be set to 0.

Start of data offset The offset, in bytes, from the start of the field data to the start of the data stream. Typically, this immediately follows the start of scan data.

Note

The last two fields have been added since the original Motion-JPEG specification, and so they may be missing from some Motion-JPEG A files. You should check the length of the APP₁ marker before using the Start Of Scan offset and Start of data offset fields. ♦

Motion-JPEG format B does not support markers. In place of the marker, therefore, QuickTime inserts a header at the beginning of the bitstream. Again, all of the fields are 32-bit integers.

Field descriptions

Reserved	Unpredictable; should be set to 0.
Tag	The data type; this field must be set to 'mjpg'.
Field size	The actual size of the image data for this field, in bytes.
Padded field size	The size of the image data, including pad bytes. Some video hardware may append pad bytes to the image data; this field, along with the field size field, allows you to compute how many pad bytes were added.
Offset to next field	The offset, in bytes, from the start of the field data to the start of the next field in the bitstream. This field should be set to 0 in the second field's header data.
Quantization Table offset	The offset, in bytes, from the start of the field data to the quantization table. If this field is set to 0, check the image description for a default quantization table.
Huffman Table offset	The offset, in bytes, from the start of the field data to the Huffman table. If this field is set to 0, check the image description for a default Huffman table.

Media Data Atom Types

Start Of Frame offset

The offset from the start of the field data to the field's image data. This field should never be set to 0.

Start Of Scan offset

The offset, in bytes, from the start of the field data to the start of scan data.

Start of data offset The offset, in bytes, from the start of the field data to the start of the data stream. Typically, this immediately follows the start of scan data.

Note

The last two fields were “reserved, must be set to zero” in the original Motion-JPEG specification. ♦

The Motion-JPEG format B header must be a multiple of 16 in size. When you add pad bytes to the header, set them to 0.

Because Motion-JPEG format B does not support markers, the JPEG bitstream does not have null bytes (0x00) inserted after data bytes that are set to 0xFF.

Figure 3-1 Motion-JPEG A dual-field sample data





Figure 3-2 Motion-JPEG B dual-field sample data



Sound Media

Sound media is used to store compressed and uncompressed audio data in QuickTime movies. It has a media type of 'soun'.

Sound Sample Description

The sound sample description contains information that defines how to interpret sound media data. This sample description is based on the standard sample description, as described in “Sample Description Atoms” (page 87).

The data format field contains the format of the audio data. Table 3-3 shows a list of supported sound codec formats.

The sound media handler also adds some of its own fields to the sample description.

Field descriptions

Version	A 16-bit integer that must be set to 0 or 1 (see below).
Revision level	A 16-bit integer that must be set to 0.
Vendor	A 32-bit integer that must be set to 0.
Number of channels	A 16-bit integer that indicates the number of sound channels used by the sound sample. Set this field to 1 for monaural sounds; set it to 2 for stereo sounds.
Sample size	A 16-bit integer that specifies the number of bits in each uncompressed sound sample. Set this field to 8 for 8-bit sound, and to 16 for 16-bit sound.
Compression ID	A 16-bit integer that must be set to 0 (or, for version 1, may be -2. See below).
Packet size	A 16-bit integer that must be set to 0.
Sample rate	A 32-bit unsigned fixed-point number that indicates the rate at which the sound samples were obtained. This number should match the media’s time scale, that is, the integer portion should match.

Variants of the Sample Description

A QuickTime sound sample description describes the format of a collection of audio samples. There are, however, three variants of the sound sample description.

Initially, QuickTime defined a sample of sound as an *uncompressed* sample, not, as in other tracks, a compressed sample. In any codec that operated on frames of audio (for example, the 20 millisecond frame common in telephony), it was necessary to determine how many uncompressed samples were in a frame, how many bytes that compressed to, and the size of the resulting uncompressed samples. This information was classically retrieved from the sound decompressor component. The sample tables (sample to time, sample size) were very compact, as they documented uncompressed samples.

Version 0 of the sound description format assumes that this is the case.

In version 1 of the sound description, introduced in QuickTime 3, the record is extended by 4 fields, each 4 bytes long, storing data that could otherwise have only been retrieved from the decompressor. This enables operating on the file in the absence of the decompressor component.

These added fields are used to support compression algorithms that can be run at different compression ratios, and to support more generic parsing of QuickTime sound tracks, that is, it allows the parsing without requiring the particular to be available.

The version field is set to 1 to indicate this new format. The four fields are:

- samples per packet—the number of uncompressed samples in a packet
- bytes per packet—the resulting compressed number of bytes for one channel
- bytes per frame—the resulting compressed number of bytes for all channels (channels * bytes per packet)
- bytes per sample—the size of an uncompressed sample

In both formats, version 0 and version 1, the timescale of the track is set to the sampling rate.

In many respects, version 1 of the `SoundDescription` record is a superset of the version 0 sound description. The new fields are taken directly from the `CompressionInfo` structure used by the Sound Manager (which uses 16-bit values) to describe the compression ratio of fixed ratio audio compression algorithms. If these fields are not used, they are set to 0. File readers only need to check to see if `samplesPerPacket` is 0.

Version 1 also defines how extensions are added to the `SoundDescription` record.

```
struct SoundDescriptionV1 {
    // original fields
    SoundDescription desc;
    // fixed compression ratio information
    unsigned long samplesPerPacket;
    unsigned long bytesPerPacket;
    unsigned long bytesPerFrame;
    unsigned long bytesPerSample;
    // optional, additional atom-based fields --
    // ([long size, long type, some data], repeat)
};
```

All other additions to the `SoundDescription` record are made using QT atoms. That means one or more atoms can be appended to the end of the `SoundDescription` record using the standard [size, type] mechanism used throughout the QuickTime movie resource architecture.

Two extensions are defined to the `SoundDescription` record. The second extension provides the ability to store data specific to a given audio decompressor in the `SoundDescription` record. Some audio decompression algorithms, such as Microsoft's ADPCM, require a set of out-of-band values to configure the decompressor. These are stored in an atom of type `siDecompressorSettings`. The contents of the `siDecompressorSettings` atom are dependent on the audio decompressor. If the QuickTime movie is created from a WAVE (.wav) or AVI (.avi) file, the `siDecompressorSettings` atom is automatically created and set to the contents of the `WAVEFORMATEX` structure from that file. In this case, the `siDecompressorSettings` atom contains little-endian data.

At runtime, the contents of the type `siSlopeAndIntercept` and `siDecompressorSettings` atoms are provided to the decompressor component through the standard `SetInfo` mechanism of the Sound Manager.

The first extension to the `SoundDescription` record is the `slope`, `intercept`, `minClip`, and `maxClip` parameters for audio. This is represented as an atom of type `siSlopeAndIntercept`. The contents of the atom are

```
struct SoundSlopeAndInterceptRecord {
    Float64 slope;
    Float64 intercept;
    Float64 minClip;
```

Media Data Atom Types

```
Float64                maxClip;
};
typedef struct SoundSlopeAndInterceptRecord SoundSlopeAndInterceptRecord;
```

VBR Audio (the Third Variant)

With the introduction of support for the playback of **variable bit-rate** (VBR) audio in QuickTime 4.1, the meaning of a number of fields has changed. (Previous versions of QuickTime supported only **constant bit-rate** (CBR) audio, while a number of modern audio compression formats—such as MP3—either support or require VBR encoding.)

In particular, a sample in a VBR audio track is, in QuickTime 4.1, a compressed frame of audio. The sample table documents these frames; specifically, the sample size table documents the size of the frames. This is constant for CBR audio, but varies for VBR.

The time to sample table documents the duration of the frames. Since the timescale is set to the sampling rate, this indicates the number of uncompressed samples in each packet, which is usually constant, even for VBR (it is common to use a fixed frame duration), though not required.

To indicate that this new meaning is used, a version 1 sound description is used and the compression ID field is set to -2. The samples per packet should be set correctly, if it is constant, as should the bytes per sample; the other two new fields are reserved and should be set to 0.

Sound Sample Data

The format of data stored in sound samples is completely dependent on the type of the compressed data stored in the sound sample description. The following sections discuss each of the formats supported by QuickTime.

Sound Codec Formats Supported

The sound codec formats shown in Table 3-3 are typical of the different forms of uncompressed audio supported in QuickTime. Note that the 16-bit forms come in both big- and little-endian forms.

Table 3-3 Supported QuickTime audio formats. Note that not all of these are uncompressed.

Format	4-Character code	Description
kSoundNotCompressed	'NONE'	Sound is not compressed
k8BitOffsetBinaryFormat	'raw '	Samples are stored uncompressed, in offset-binary format (values range from 0 to 255; 128 is silence). These are stored as 8-bit offset binaries.
k16BitBigEndianFormat	'twos'	Samples are stored uncompressed, in two's-complement format (sample values range from -128 to 127 for 8-bit audio, and -32768 to 32767 for 16-bit audio; 0 is always silence). These samples are stored in 16-bit big-endian format.
k16BitLittleEndianFormat	'sowt'	16-bit little-endian, twos-complement
kMACE3Compression	'MAC3'	Samples have been compressed using MACE 3:1.
kMACE6Compression	'MAC6'	Samples have been compressed using MACE 6:1.
kIMACompression	'ima4'	Samples have been compressed using IMA 4:1.
kFloat32Format	'fl32'	32-bit floating point
kFloat64Format	'fl64'	64-bit floating point
k24BitFormat	'in24'	24-bit integer
k32BitFormat	'in32'	32-bit integer
kULawCompression	'ulaw'	uLaw 2:1
kALawCompression	'alaw'	uLaw 2:1
kMicrosoftADPCMFormat	0x6D730002	Microsoft ADPCM-ACM code 2
kDVIIntelIMAFormat	0x6D730011	DVI/Intel IMAADPCM-ACM code 17

Media Data Atom Types

Format	4-Character code	Description
kDVAudioFormat	'dvca'	DV Audio
kQDesignCompression	'QDMC'	QDesign music
unamed	'QDM2'	QDesign music version 2 (no constant).
kQUALCOMMCompression	'Qc1p'	QUALCOMM PureVoice
kMPGELayer3Format	0x6D730055	MPEG layer 3, CBR only (pre-QT4.1)
kFullMPEGLayer3Format	' .mp3'	MPEG layer 3, CBR & VBR (QT4.1 and later)

Uncompressed 8-Bit Sound

Eight-bit audio is stored in offset-binary encodings. If the data is in stereo, the left and right channels are interleaved.

Uncompressed 16-Bit Sound

Sixteen-bit audio may be stored in two's-complement encodings. If the data is in stereo, the left and right channels are interleaved.

Formats Not Currently in Use: MACE 3:1 and 6:1

These compression formats are obsolete: MACE 3:1 and 6:1

These are 8-bit sound codec formats, defined as follows:

```
kMACE3Compression = FOUR_CHAR_CODE('MAC3'), /*MACE 3:1*/
kMACE6Compression = FOUR_CHAR_CODE('MAC6'), /*MACE 6:1*/
```

IMA, uLaw, and aLaw

■ IMA 4:1

The IMA encoding scheme is based on a standard developed by the International Multimedia Association for pulse code modulation (PCM) audio compression. QuickTime uses a slight variation of the format to allow for

Media Data Atom Types

random access. IMA is a 16-bit audio format which supports 4:1 compression. It is defined as follows:

```
kIMACompression = FOUR_CHAR_CODE('ima4'), /*IMA 4:1*/
```

■ uLaw 2:1 and aLaw 2:1

The uLaw (mu-law) encoding scheme is used on North American and Japanese phone systems, and is coming into use for voice data interchange, and in PBXs, voice-mail systems, and Internet talk radio (via MIME). In uLaw encoding, 14 bits of linear sample data are reduced to 8 bits of logarithmic data.

The aLaw encoding scheme is used in Europe and the rest of the world.

The `kULawCompression` and the `kALawCompression` formats are typically found in `.au` formats.

Floating-Point Formats

Both `kFloat32Format` and `kFloat64Format` are floating-point uncompressed formats. Depending upon codec-specific data associated with the sample description, the floating-point values may be in big-endian (network) or little-endian (Intel) byte order. This differs from the 16-bit formats, where there is a single format for each endia layout.

24- and 32-Bit Integer Formats

Both `k24BitFormat` and `k32BitFormat` are integer uncompressed formats. Depending upon codec-specific data associated with the sample description, the floating-point values may be in big-endian (network) or little-endian (Intel) byte order.

kMicrosoftADPCMFormat and kDVIIntelIMAFormat Sound Codecs

The `kMicrosoftADPCMFormat` and the `kDVIIntelIMAFormat` codec provide QuickTime interoperability with AVI and WAV files. The four-character codes used by Microsoft for their formats are numeric. To construct a QuickTime-supported codec format of this type, the Microsoft numeric ID is taken to generate a four-character code of the form `'msxx'` where `xx` takes on the numeric ID.

kDVAudioFormat Sound Codec

The DV audio sound codec, `kDVAudioFormat`, decodes audio found in a DV stream. Since a DV frame contains both video and audio, this codec knows how to skip video portions of the frame and only retrieve the audio portions. Likewise, the video codec skips the audio portions and renders only the image.

kQDesignCompression Sound Codec

The `kQDesignCompression` sound codec is the QDesign 1 (pre-QuickTime 4) format. Note that there is also a QDesign 2 format whose four-character code is 'QDM2'.

MPEG Layer 3 (MP3) Codecs

The QuickTime MPEG layer 3 (MP3) codecs come in two particular flavors, as shown in Table 3-3. The first (`kMPEGLayer3Format`) is used exclusively in the constant bitrate (CBR) case (pre-QuickTime 4). The other (`kFullMPEGLay3Format`) is used in both the CBR and variable bitrate (VBR) cases. Note that they are the same codec underneath.

Timecode Media

Timecode media is used to store time code data in QuickTime movies. It has a media type of 'tmcd'.

Timecode Sample Description

The timecode sample description contains information that defines how to interpret time code media data. This sample description is based on the standard sample description header, as described in “Sample Description Atoms” (page 87).

The data format field in the sample description is always set to 'tmcd'.

The timecode media handler also adds some of its own fields to the sample description.

Field descriptions

Reserved	A 32-bit integer that is reserved for future use. Set this field to 0.
Flags	A 32-bit integer containing flags that identify some timecode characteristics. The following flags are defined.
Drop frame	Indicates whether the timecode is drop frame. Set it to 1 if the timecode is drop frame. This flag's value is 0x0001.
24 hour max	Indicates whether the timecode wraps after 24 hours. Set it to 1 if the timecode wraps. This flag's value is 0x0002.
Negative times OK	Indicates whether negative time values are allowed. Set it to 1 if the timecode supports negative values. This flag's value is 0x0004.
Counter	Indicates whether the time value corresponds to a tape counter value. Set it to 1 if the timecode values are tape counter values. This flag's value is 0x0008.
Time scale	A 32-bit integer that specifies the time scale for interpreting the frame duration field.
Frame duration	A 32-bit integer that indicates how long each frame lasts in real time.
Number of frames	An 8-bit integer that contains the number of frames per second for the timecode format. If the time is a counter, this is the number of frames for each counter tick.
Reserved	A 24-bit quantity that must be set to 0.
Source reference	A user data atom containing information about the source tape. The only currently used user data list entry is the 'name' type. This entry contains a text item specifying the name of the source tape.

Timecode Media Information Atom

The timecode media also requires a media information atom. This atom contains information governing how the timecode text is displayed. This media information atom is stored in a base media information atom (see “Base Media Information Atoms” (page 77) for more information). The type of the timecode media information atom is 'tcmi'.

The timecode media information atom contains the following data elements.

Field descriptions

Size	A 32-bit integer that specifies the number of bytes in this time code media information atom.
Type	A 32-bit integer that identifies the atom type; this field must be set to 'tcmi'.
Version	A 1-byte specification of the version of this timecode media information atom.
Flags	A 3-byte space for timecode media information flags. Set this field to 0.
Text font	A 16-bit integer that indicates the font to use. Set this field to 0 to use the system font. If the font name field contains a valid name, ignore this field.
Text face	A 16-bit integer that indicates the font's style. Set this field to 0 for normal text. You can enable other style options by using one or more of the following bit masks: 0x0001 Bold 0x0002 Italic 0x0004 Underline 0x0008 Outline 0x0010 Shadow 0x0020 Condense 0x0040 Extend
Text size	A 16-bit integer that specifies the point size of the time code text.
Text color	A 48-bit RGB color value for the timecode text.
Background color	A 48-bit RGB background color for the timecode text.
Font name	A Pascal string specifying the name of the timecode text's font.

Timecode Sample Data

There are two different sample data formats used by timecode media.

If the Counter flag is set to 1 in the timecode sample description, the sample data is a counter value. Each sample contains a 32-bit integer counter value.

If the Counter flag is set to 0 in the timecode sample description, the sample data format is a timecode record, as follows.

Field descriptions

Hours	An 8-bit unsigned integer that indicates the starting number of hours.
Negative	A 1-bit value indicating the time's sign. If bit is set to 1, the timecode record value is negative.
Minutes	A 7-bit integer that contains the starting number of minutes.
Seconds	An 8-bit unsigned integer indicating the starting number of seconds.
Frames	An 8-bit unsigned integer that specifies the starting number of frames. This field's value cannot exceed the value of the number of frames field in the timecode sample description.

Text Media

Text media is used to store text data in QuickTime movies. It has a media type of 'text'.

Text Sample Description

The text sample description contains information that defines how to interpret text media data. This sample description is based on the standard sample description header, as described in “Sample Description Atoms” (page 87).

The data format field in the sample description is always set to 'text'.

The text media handler also adds some of its own fields to the sample description.

Field descriptions

Display flags A 32-bit integer containing flags that describe how the text should be drawn. The following flags are defined.

Don't auto scale

Controls text scaling. If this flag is set to 1, the text media handler reflows the text instead of scaling when the track is scaled. This flag's value is 0x0002.

Use movie background color

Controls background color. If this flag is set to 1, the text media handler ignores the background color field in the text sample description and uses the movie's background color instead. This flag's value is 0x0008.

Scroll in

Controls text scrolling. If this flag is set to 1, the text media handler scrolls the text until the last of the text is in view. This flag's value is 0x0020.

Scroll out

Controls text scrolling. If this flag is set to 1, the text media handler scrolls the text until the last of the text is gone. This flag's value is 0x0040.

Horizontal scroll

Controls text scrolling. If this flag is set to 1, the text media handler scrolls the text horizontally; otherwise, it scrolls the text vertically. This flag's value is 0x0080.

Reverse scroll

Controls text scrolling. If this flag is set to 1, the text media handler scrolls down (if scrolling vertically) or backward (if scrolling horizontally; note that horizontal scrolling also depends upon text justification). This flag's value is 0x0100.

Continuous scroll

Controls text scrolling. If this flag is set to 1, the text media handler displays new samples by scrolling out the old ones. This flag's value is 0x0200.

Drop shadow

Controls drop shadow. If this flag is set to 1, the text media handler displays the text with a drop shadow. This flag's value is 0x1000.

Anti-alias

Controls anti-aliasing. If this flag is set to 1, the text media handler uses anti-aliasing when drawing text. This flag's value is 0x2000.

Key text

Controls background color. If this flag is set to 1, the text media handler does not display the background color, so that the text overlay background tracks. This flag's value is 0x4000.

Text justification	A 32-bit integer that indicates how the text should be aligned. Set this field to 0 for left-justified text, to 1 for centered text, and to -1 for right-justified text.
Background color	A 48-bit RGB color that specifies the text's background color.
Default text box	A 64-bit rectangle that specifies an area to receive text (top, left, bottom, right). Typically this field is set to all zeros.
Reserved	A 64-bit value that must be set to 0.
Font number	A 16-bit value that must be set to 0.
Font face	A 16-bit integer that indicates the font's style. Set this field to 0 for normal text. You can enable other style options by using one or more of the following bit masks: 0x0001 Bold 0x0002 Italic 0x0004 Underline 0x0008 Outline 0x0010 Shadow 0x0020 Condense

Media Data Atom Types

	0x0040 Extend
Reserved	An 8-bit value that must be set to 0.
Reserved	A 16-bit value that must be set to 0.
Foreground color	A 48-bit RGB color that specifies the text's foreground color.
Text name	A Pascal string specifying the name of the font to use to display the text.

Text Sample Data

The format of the text data is a 16-bit length word followed by the actual text. The length word specifies the number of bytes of text, not including the length word itself. Following the text, there may be one or more atoms containing additional information for drawing and searching the text.

Table 3-4 lists the currently defined text sample extensions.

Table 3-4 Text sample extensions

Text sample extension	Description
'styl'	Style information for the text. Allows you to override the default style in the sample description or to define more than one style for a sample. The data is a TextEdit style scrap.
'ftab'	Table of font names. Each table entry contains a font number (stored in a 16-bit integer) and a font name (stored in a Pascal string). This atom is required if the 'styl' atom is present.
'hlit'	Highlight information. The atom data consists of two 32-bit integers. The first contains the starting offset for the highlighted text, and the second has the ending offset. A highlight sample can be in a key frame or in a differenced frame. When it's used in a differenced frame, the sample should contain a zero-length piece of text.
'hclr'	Highlight color. This atom specifies the 48-bit RGB color to use for highlighting.

Table 3-4 Text sample extensions

Text sample extension	Description
'drpo'	Drop shadow offset. When the display flags indicate drop shadow style, this atom can be used to override the default drop shadow placement. The data consists of two 16-bit integers. The first indicates the horizontal displacement of the drop shadow, in pixels; the second, the vertical displacement.
'drpt'	Drop shadow transparency. The data is a 16-bit integer between 0 and 256 indicating the degree of transparency of the drop shadow. A value of 256 makes the drop shadow completely opaque.
'imag'	Image font data. This atom contains two more atoms. An 'idat' atom contains compressed image data to be used to draw the text when the required fonts are not available. An 'idsc' atom contains a video sample description describing the format of the compressed image data.
'metr'	Image font highlighting. This atom contains metric information that governs highlighting when an 'imag' atom is used for drawing.

Hypertext

Hypertext is used as an action that takes you to a Web URL; like a Web URL, it appears blue and underlined. Hypertext is stored in a text track sample atom stream as type 'htxt'.

The data stored is a `QTAtomContainer`. The root atom of hypertext in this container is of type 'wtxt'. This is the parent for all individual hypertext objects.

For each hypertext item, the parent atom is of type 'htxt'. This is the atom container atom type. Two children of this atom that define the offset of the hypertext in the text stream are

kRangeStart	'strt' // unsigned long
kRangeEnd	'end ' // unsigned long

Media Data Atom Types

Child atoms of the parent atom are the events of type `kQTEventType` and the ID of the event type. The children of these event atoms follow the same format as other wired events.

```
kTextWiredObjectsAtomType, 1
    kHyperTextItemAtomType, 1..n
        kRangeStart, 1
            long
        kRangeEnd, 1
            long

    kQTEventType, (kQTEventMouseClicked, kQTEventMouseClickedEnd,
        kQTEventMouseClickedEndTriggerButton,
        kQTEventMouseEnter, kQTEventMouseExit)
    kAction        // The known range of track movie sprite actions
...

```

Music Media

Music media is used to store note-based audio data, such as MIDI data, in QuickTime movies. It has a media type of `'musi'`.

Music Sample Description

The music sample description uses the standard sample description header, as described in the section “Sample Description Atoms” (page 87).

The data format field in the sample description is always set to `'musi'`. The music media handler adds an additional 32-bit integer field to the sample description containing flags. Currently no flags are defined, and this field should be set to 0.

Following the flags field, there may be appended data in the QuickTime music format. This data consists of part-to-instrument mappings in the form of General events containing note requests. One note request event should be present for each part that will be used in the sample data.

Music Sample Data

The sample data for music samples consists entirely of data in the QuickTime music format. Typically, up to 30 seconds of notes are grouped into a single sample. For a complete description of the QuickTime music format, see your most recent QuickTime developer documentation.

MPEG Media

MPEG media is used to store MPEG streams in QuickTime movies. It has a media type of 'MPEG'.

MPEG Sample Description

The MPEG sample description uses the standard sample description header, as described in “Sample Description Atoms” (page 87).

The data format field in the sample description is always set to 'MPEG'. The MPEG media handler adds no additional fields to the sample description.

MPEG Sample Data

Each sample in an MPEG media is an entire MPEG stream. This means that a single MPEG sample may be several hundred megabytes in size. The MPEG encoding used by QuickTime corresponds to the ISO standard, as described in ISO document CD 11172.

Sprite Media

Sprite media is used to store character-based animation data in QuickTime movies. It has a media type of 'sprt'.

Sprite Sample Description

The sprite sample description uses the standard sample description header, as described in “Sample Description Atoms” (page 87).

The data format field in the sample description is always set to 'sprt'. The sprite media handler adds no additional fields to the sample description.

Sprite Sample Data

All sprite samples are stored in QT atom structures. The sprite media uses both key frames and differenced frames. The key frames contain all of the sprite's image data, and the initial settings for each of the sprite's properties.

A key frame always contains a shared data atom of type 'dflt'. This atom contains data to be shared between the sprites, consisting mainly of image data and sample descriptions. The shared data atom contains a single sprite image container atom, with an atom type value of 'imct' and an ID value of 1.

The sprite image container atom stores one or more sprite image atoms of type 'imag'. Each sprite image atom contains an image sample description immediately followed by the sprite's compressed image data. The sprite image atoms should have ID numbers starting at 1 and counting consecutively upward.

The key frame also must contain definitions for each sprite in atoms of type 'sprt'. Sprite atoms should have ID numbers start at 1 and count consecutively upward. Each sprite atom contains a list of properties. Table 3-5 shows all currently defined sprite properties.

Table 3-5 Sprite properties

Property name	Value	Description
kSpritePropertyMatrix	1	Describes the sprite's location and scaling within its sprite world or sprite track. By modifying a sprite's matrix, you can modify the sprite's location so that it appears to move in a smooth path on the screen or so that it jumps from one place to another. You can modify a sprite's size, so that it shrinks, grows, or stretches. Depending on which image compressor is used to create the sprite images, other transformations, such as rotation, may be supported as well. Translation-only matrices provide the best performance.
kSpritePropertyVisible	4	Specifies whether or not the sprite is visible. To make a sprite visible, you set the sprite's visible property to true.
kSpritePropertyLayer	5	Contains a 16-bit integer value specifying the layer into which the sprite is to be drawn. Sprites with lower layer numbers appear in front of sprites with higher layer numbers. To designate a sprite as a background sprite, you should assign it the special layer number kBackgroundSpriteLayerNum.

Property name	Value	Description
<code>kSpritePropertyGraphicsMode</code>	6	Specifies a graphics mode and blend color that indicates how to blend a sprite with any sprites behind it and with the background. To set a sprite's graphics mode, you call <code>SetSpriteProperty</code> , passing a pointer to a <code>ModifierTrackGraphicsModeRecord</code> structure.
<code>kSpritePropertyActionHandlingSpriteID</code>	8	Specifies another sprite by ID that delegates QT events.
<code>kSpritePropertyImageIndex</code>	100	Contains the atom ID of the sprite's image atom.

The override sample differs from the key frame sample in two ways. First, the override sample does not contain a shared data atom. All shared data must appear in the key frame. Second, only those sprite properties that change need to be specified. If none of a sprite's properties change in a given frame, then the sprite does not need an atom in the differenced frame.

The override sample can be used in one of two ways: combined, as with video key frames, to construct the current frame; or the current frame can be derived by combining only the key frame and the current override sample.

Refer to the section “Sprite Track Media Format” (page 145) for information on how override samples are indicated in the file, using `kSpriteTrackPropertySampleFormat` and the default behavior of the `kKeyFrameAndSingleOverride` format.

Sprite Track Properties

In addition to defining properties for individual sprites, you can also define properties that apply to an entire **sprite track**. These properties may override default behavior or provide hints to the sprite media handler. The following sprite track properties are supported:

- The `kSpriteTrackPropertyBackgroundColor` property specifies a background color for the sprite track. The background color is used for any area that is not covered by regular sprites or background sprites. If you do not specify a background color, the sprite track uses black as the default background color.

Media Data Atom Types

- The `kSpriteTrackPropertyOffscreenBitDepth` property specifies a preferred bit depth for the sprite track's offscreen buffer. The allowable values are 8 and 16. To save memory, you should set the value of this property to the minimum depth needed. If you do not specify a bit depth, the sprite track allocates an offscreen buffer with the depth of the deepest intersecting monitor.
- The `kSpriteTrackPropertySampleFormat` property specifies the sample format for the sprite track. If you do not specify a sample format, the sprite track uses the default format, `kKeyFrameAndSingleOverride`.

To specify sprite track properties, you create a single QT atom container and add a leaf atom for each property you want to specify. To add the properties to a sprite track, you call the media handler function `SetMediaPropertyAtom`. To retrieve a sprite track's properties, you call the media handler function `GetMediaPropertyAtom`.

The sprite track properties and their corresponding data types are listed in Table 3-6.

Table 3-6 Sprite track properties

Atom type	Atom ID	Leaf data type
<code>kSpriteTrackPropertyBackgroundColor</code>	1	<code>RGBColor</code>
<code>kSpriteTrackPropertyOffscreenBitDepth</code>	1	unsigned short
<code>kSpriteTrackPropertySampleFormat</code>	1	long
<code>kSpriteTrackPropertyHasActions</code>	1	Boolean
<code>kSpriteTrackPropertyQTIdleEventsFrequency</code>	1	<code>UInt32</code>
<code>kSpriteTrackPropertyVisible</code>	1	Boolean
<code>kSpriteTrackPropertyScaleSpritesToScaleWorld</code>	1	Boolean

Note


When pasting portions of two different tracks together, the Movie Toolbox checks to see that all sprite track properties match. If, in fact, they do match, the paste results in a single sprite track instead of two. ♦

Sprite Track Media Format

The sprite track media format is hierarchical and based on QT atoms and atom containers. A sprite track is defined by one or more key frame samples, each followed by any number of override samples. A key frame sample and its subsequent override samples define a scene in the sprite track. A key frame sample is a QT atom container that contains atoms defining the sprites in the scene and their initial properties. The override samples are other QT atom containers that contain atoms that modify sprite properties, thereby animating the sprites in the scene. In addition to defining properties for individual sprites, you can also define properties that apply to an entire sprite track.

Figure 3-3 shows the high-level structure of a sprite track key frame sample. Each atom in the atom container is represented by its atom type, atom ID, and, if it is a leaf atom, the type of its data.


Figure 3-3 A key frame sample atom container



The QT atom container contains one child atom for each sprite in the key frame sample. Each sprite atom has a type of `kSpriteAtomType`. The sprite IDs are numbered from 1 to the number of sprites defined by the key frame sample (`numSprites`).


Each sprite atom contains leaf atoms that define the properties of the sprite, as shown in Figure 3-4. For example, the `kSpritePropertyLayer` property defines a sprite's layer. Each sprite property atom has an atom type that corresponds to the property and an ID of 1.

Figure 3-4 Atoms that describe a sprite and its properties



In addition to the sprite atoms, the QT atom container contains one atom of type `kSpriteSharedDataAtomType` with an ID of 1. The atoms contained by the shared data atom describe data that is shared by all sprites. The shared data atom contains one atom of type `kSpriteImagesContainerAtomType` with an ID of 1 (Figure 3-5).

The image container atom contains one atom of type `kImageAtomType` for each image in the key frame sample. The image atom IDs are numbered from 1 to the number of images (`numImages`). Each image atom contains a leaf atom that holds the image data (type `kSpriteImageDataAtomType`) and an optional leaf atom (type `kSpriteNameAtomType`) that holds the name of the image.

Figure 3-5 Atoms that describe sprite images

Sprite Media Format Atoms

The sprite track's sample format enables you to store the atoms necessary to describe action lists that are executed in response to QuickTime events. "QT Atom Container Description Key" (page 156) defines a grammar for constructing valid action sprite samples, which may include complex expressions.

Both key frame samples and override samples support the sprite action atoms. Override samples override actions at the QuickTime event level. In effect, what you do by overriding is to completely replace one event handler and all its actions with another. The sprite track's `kSpriteTrackPropertySampleFormat` property has no effect on how actions are performed. The behavior is similar to the default `kKeyFrameAndSingleOverride` format where, if in a given override sample there is no handler for the event, the key frame's handler is used, if there is one.

Sprite Media Format Extensions

This section describes some of the atom types and IDs used to extend the sprite track's media format, thus enabling action sprite capabilities.

A complete description of the grammar for sprite media handler samples, including action sprite extensions, is included in the section “Sprite Media Handler Track Properties QT Atom Container Format” (page 157).

IMPORTANT

Some sprite track property atoms were added in QuickTime 4. In particular, you must set the `kSpriteTrackPropertyHasActions` track property in order for your sprite actions to be executed. ▲

Sprite Track Property Atoms

Sprite track property atoms are applied to the whole track, not just to a single sample.

Constant descriptions

`kSpriteTrackPropertyHasActions`

You must add an atom of this type with its leaf data set to `true` if you want the movie controller to execute the actions in your sprite track's media. The atom's leaf data is of type `Boolean`. The default value is `false`, so it is very important to add an atom of this type if you want interactivity to take place.

`kSpriteTrackPropertyQTIdleEventsFrequency`

You must add an atom of this type if you want the sprites in your sprite track to receive `kQTEventIdle` QuickTime events. The atom's leaf data is of type `UInt32`. The value is the minimum number of ticks that must pass before the next `QTIdle` event is sent. Each tick is 1/60th of one second. To specify “Idle as fast as possible,” set the value to 0. The default value is `kNoQTIdleEvents`, which means don't send any idle events.

It is possible that for small idle event frequencies, the movie will not be able to keep up, in which case idle events will be sent as fast as possible.

Media Data Atom Types

Since sending idle events takes up some time, it is best to specify the largest frequency that produces the results that you desire, or `kNoQTIdleEvents` if you do not need them.

`kSpriteTrackPropertyVisible`

You can cause the entire sprite track to be invisible by setting the value of this `Boolean` property to `false`. This is useful for using a sprite track as a hidden button track—for example, placing an invisible sprite track over a video track would allow the characters in the video to be clickable. The default value is `visible` (`true`).

`kSpriteTrackPropertyScaleSpritesToScaleWorld`

You can cause each sprite to be rescaled when the sprite track is resized by setting the value of this `Boolean` property to `true`. Setting this property can improve the drawing performance and quality of a scaled sprite track. This is particularly useful for sprite images compressed with codecs that are resolution-independent, such as the Curve codec. The default value for this property is `false`.

Atom Types

The following constants represent atom types for sprite media.

```
enum {
    kSpriteAtomType           = 'sprt',
    kSpriteImagesContainerAtomType = 'imct',
    kSpriteImageAtomType      = 'imag',
    kSpriteImageDataAtomType  = 'imda',
    kSpriteImageDataRefAtomType = 'imre',
    kSpriteImageDataRefTypeAtomType = 'imrt',
    kSpriteImageGroupIDAtomType = 'imgr',
    kSpriteImageRegistrationAtomType = 'imrg',
    kSpriteImageDefaultImageIndexAtomType = 'defi',
    kSpriteSharedDataAtomType = 'dflt',
    kSpriteNameAtomType      = 'name',
    kSpriteImageNameAtomType = 'name',
    kSpriteUsesImageIDsAtomType = 'uses',
    kSpriteBehaviorsAtomType = 'beha',
    kSpriteImageBehaviorAtomType = 'imag',
    kSpriteCursorBehaviorAtomType = 'crsr',
```

Media Data Atom Types

```

kSpriteStatusStringsBehaviorAtomType = 'sstr',
kSpriteVariablesContainerAtomType    = 'vars',
kSpriteStringVariableAtomType        = 'strv',
kSpriteFloatingPointVariableAtomType = 'flow'
kSpriteSharedDataAtomType             = 'dflt',
kSpriteURLLinkAtomType                = 'url '
kSpritePropertyMatrix                 = 1
kSpritePropertyVisible                = 4
kSpritePropertyLayer                 = 5
kSpritePropertyGraphicsMode          = 6
kSpritePropertyImageIndex            = 100
kSpritePropertyBackgroundColor        = 101
kSpritePropertyOffscreenBitDepth     = 102
kSpritePropertySampleFormat          = 103
};

```

Constant descriptions

`kSpriteAtomType` **The atom is a parent atom that describes a sprite. It contains atoms that describe properties of the sprite. Optionally, it may also include an atom of type `kSpriteNameAtomType` that defines the name of the sprite.**

`kSpriteImagesContainerAtomType` **The atom is a parent atom that contains atoms of type `kSpriteImageAtomType`.**

`kSpriteImageAtomType` **The atom is a parent atom that contains an atom of type `kSpriteImageDataAtomType`. Optionally, it may also include an atom of type `kSpriteNameAtomType` that defines the name of the image.**

`kSpriteImageDataAtomType` **The atom is a leaf atom that contains image data.**

`kSpriteSharedDataAtomType` **The atom is a parent atom that contains shared sprite data, such as an atom container of type `kSpriteImagesContainerAtomType`.**

`kSpriteNameAtomType` **The atom is a leaf atom that contains the name of a sprite or an image. The leaf data is composed of one or more ASCII characters.**

Media Data Atom Types

`kSpritePropertyImageIndex`

A leaf atom containing the image index property which is of type `short`. This atom is a child atom of `kSpriteAtom`.

`kSpritePropertyLayer`

A leaf atom containing the layer property which is of type `short`. This atom is a child atom of `kSpriteAtom`.

`kSpritePropertyMatrix`

A leaf atom containing the matrix property which is of type `MatrixRecord`. This atom is a child atom of `kSpriteAtom`.

`kSpritePropertyVisible`

A leaf atom containing the visible property which is of type `short`. This atom is a child atom of `kSpriteAtom`.

`kSpritePropertyGraphicsMode`

A leaf atom containing the graphics mode property which is of type `ModifierTrackGraphicsModeRecord`. This atom is a child atom of `kSpriteAtom`.

`kSpritePropertyBackgroundColor`

A leaf atom containing the background color property which is of type `RGBColor`. This atom is used in a sprite track's `MediaPropertyAtom` atom container.

`kSpritePropertyOffscreenBitDepth`

A leaf atom containing the preferred offscreen bitdepth which is of type `short`. This atom is used in a sprite track's `MediaPropertyAtom` atom container.

`kSpritePropertySampleFormat`

A leaf atom containing the sample format property, which is of type `short`. This atom is used in a sprite track's `MediaPropertyAtom` atom container.

`kSpriteImageRegistrationAtomType`

Sprite images have a default registration point of 0, 0. To specify a different point, add an atom of type `kSpriteImageRegistrationAtomType` as a child atom of the `kSpriteImageAtomType` and set its leaf data to a `FixedPoint` value with the desired registration point.

`kSpriteImageGroupIDAtomType`

You must assign group IDs to sets of equivalent images in your key frame sample. For example, if the sample contains ten images where the first two images are equivalent, and

Media Data Atom Types

the last eight images are equivalent, then you could assign a group ID of 1000 to the first two images, and a group ID of 1001 to the last eight images. This divides the images in the sample into two sets. The actual ID does not matter, it just needs to be a unique positive integer.

Each image in a sprite media key frame sample is assigned to a group. Add an atom of type `kSpriteImageGroupIDAtomType` as a child of the `kSpriteImageAtomType` atom and set its leaf data to a long containing the group ID.

IMPORTANT

You must assign group IDs to your sprite sample if you want a sprite to display images with non-equivalent image descriptions (i.e., images with different dimensions). ▲

For each of the following atom types (added to QuickTime 4)—except `kSpriteBehaviorsAtomType`—you fill in the structure `QTSpriteButtonBehaviorStruct`, which contains a value for each of the four states.

`kSpriteBehaviorsAtomType`

This is the parent atom of `kSpriteImageBehaviorAtomType`, `kSpriteCursorBehaviorAtomType`, and `kSpriteStatusStringsBehaviorAtomType`.

`kSpriteImageBehaviorAtomType`

Specifies the `imageIndex`.

`kSpriteCursorBehaviorAtomType`

Specifies the `cursorID`.

`kSpriteStatusStringsBehaviorAtomType`

Specifies an ID of a string variable contained in a sprite track to display in the status area of the browser.

Note

All sprite media—specifically the leaf data in the QT atom containers for sample and sprite track properties—should be written in big-endian format. ♦

`kSpriteUsesImageIDsAtomType`

Media Data Atom Types

This atom allows a sprite to specify which images it uses—in other words, the subset of images that its `imageIndex` property can refer to.

You add an atom of type `kSpriteUsesImageIDsAtomType` as a child of a `kSpriteAtomType` atom, setting its leaf data to an array of QT atom IDs. This array contains the IDs of the images used, not the indices.

Although QuickTime does not currently use this atom internally, tools that edit sprite media can use the information provided to optimize certain operations, such as cut, copy, and paste.

`kSpriteImageRegistrationAtomType`

Sprite images have a default registration point of 0, 0. To specify a different point, you add an atom of type `kSpriteImageRegistrationAtomType` as a child atom of the `kSpriteImageAtomType` and set its leaf data to a `FixedPoint` value with the desired registration point.

`kSpriteImageGroupIDAtomType`

You must assign group IDs to sets of equivalent images in your key frame sample. For example, if the sample contains ten images where the first two images are equivalent, and the last eight images are equivalent, then you could assign a group ID of 1000 to the first two images, and a group ID of 1001 to the last eight images. This divides the images in the sample into two sets. The actual ID does not matter; it just needs to be a unique positive integer.

Each image in a sprite media key frame sample is assigned to a group. You add an atom of type

`kSpriteImageGroupIDAtomType` as a child of the `kSpriteImageAtomType` atom and set its leaf data to a long containing the group ID.

IMPORTANT

You must assign group IDs to your sprite sample if you want a sprite to display images with non-equivalent image descriptions (that is, images with different dimensions). ▲

You use the following atom types, which were added to QuickTime 4, to specify that an image is referenced and how to access it.

`kSpriteImageDataRefAtomType`

Media Data Atom Types

Add this atom as a child of the `kSpriteImageAtomType` atom instead of a `kSpriteImageDataAtomType`. Its ID should be 1. Its data should contain the data reference (similar to the `dataRef` parameter of `GetDataHandler`).

`kSpriteImageDataRefTypeAtomType`

Add this atom as a child of the `kSpriteImageAtomType` atom. Its ID should be 1. Its data should contain the data reference type (similar to the `dataRefType` parameter of `GetDataHandler`).

`kSpriteImageDefaultImageIndexAtomType`

You may optionally add this atom as a child of the `kSpriteImageAtomType` atom. Its ID should be 1. Its data should contain a short, which specifies an image index of a traditional image to use while waiting for the referenced image to load.

The following constants represent formats of a sprite track. The value of the constant indicates how override samples in a sprite track should be interpreted. You set a sprite track's format by creating a `kSpriteTrackPropertySampleFormat` atom.

```
enum {
    kKeyFrameAndSingleOverride    = 1L << 1,
    kKeyFrameAndAllOverrides      = 1L << 2
};
```

Constant descriptions

`kKeyFrameAndSingleOverride`

The current state of the sprite track is defined by the most recent key frame sample and the current override sample. This is the default format.

`kKeyFrameAndAllOverrides`

The current state of the sprite track is defined by the most recent key frame sample and all subsequent override samples up to and including the current override sample.

Sprite Button Behaviors

In QuickTime 4, sprites in a sprite track can specify simple button behaviors. These behaviors can control the sprite's image, the system cursor, and the status message displayed in a Web browser. They also provide a shortcut for a common set of actions that may result in more efficient QuickTime movies.

Button behaviors can be added to a sprite. These behaviors are intended to make the common task of creating buttons in a sprite track easy—you basically just fill in a template.

Three types of behaviors are available; you may choose one or more behaviors. Each change a type of property associated with a button and are triggered by the mouse states `notOverNotPressed`, `overNotPressed`, `overPressed`, and `notOverPressed`. The three properties changed are

- the sprite's `imageIndex`
- the ID of a cursor to be displayed
- the ID of a status string variable displayed in the URL status area of a Web browser.

Setting a property's value to `-1` means don't change it.

Note

The cursor is automatically set back to the default system cursor when leaving a sprite. ♦

The sprite track handles letting one sprite act as an active button at a time.

The behaviors are added at the beginning of the sprite's list of actions, so they may be overridden by actions if desired.

To use the behaviors, you fill in the new atoms as follows, using the description key specified in “QT Atom Container Description Key” (page 156):

```
kSpriteAtomType
  <kSpriteBehaviorsAtomType>, 1
    <kSpriteImageBehaviorAtomType>
      [QTSpriteButtonBehaviorStruct]
    <kSpriteCursorBehaviorAtomType>
      [QTSpriteButtonBehaviorStruct]
    <kSpriteStatusStringsBehaviorAtomType>
      [QTSpriteButtonBehaviorStruct]
```

QT Atom Container Description Key

Because QT atom container–based data structures are widely used in QuickTime, a description key is presented here. Its usage is illustrated in the following sections, “Sprite Media Handler Track Properties QT Atom Container Format” (page 157) and “Sprite Media Handler Sample QT Atom Container Formats” (page 157).

```
[(QTAtomFormatName)] =
    atomType_1, id, index
    data
    atomType_n, id, index
    data
```

The atoms may be required or optional:

```
<atomType>    // optional atom
atomType      // required atom
```

The atom ID may be a number if it is required to be a constant, or it may be a list of valid atom IDs, indicating that multiple atoms of this type are allowed.

```
3                // one atom with id of 3
(1..3)           // three atoms with id's of 1, 2, and 3
(1, 5, 7)        // three atoms with id's of 1, 5, and 7
(anyUniqueIDs)   // multiple atoms each with a unique id
```

The atom index may be a 1 if only one atom of this type is allowed, or it may be a range from 1 to some constant or variable.

```
1                // one atom of this type is allowed, index is always 1
(1..3)           // three atoms with indexes 1, 2, and 3
(1..numAtoms)    // numAtoms atoms with indexes of 1 to numAtoms
```

The data may be leaf data in which its data type is listed inside of brackets [], or it may be a nested tree of atoms.

```
[theDataType]    // leaf data of type theDataType
childAtoms       // a nested tree of atoms
```

Media Data Atom Types

Nested QTAtom format definitions [(AtomFormatName)] may appear in a definition.

Sprite Media Handler Track Properties QT Atom Container Format

```
[(SpriteTrackProperties)]
  <kSpriteTrackPropertyBackgroundColor, 1, 1>
    [RGBColor]
  <kSpriteTrackPropertyOffscreenBitDepth, 1, 1>
    [short]
  <kSpriteTrackPropertySampleFormat, 1, 1>
    [long]
  <kSpriteTrackPropertyScaleSpritesToScaleWorld, 1, 1>
    [Boolean]
  <kSpriteTrackPropertyHasActions, 1, 1>
    [Boolean]
  <kSpriteTrackPropertyVisible, 1, 1>
    [Boolean]
  <kSpriteTrackPropertyQTIdleEventsFrequency, 1, 1>
    [UInt32]
```

Sprite Media Handler Sample QT Atom Container Formats

```
[(SpriteKeySample)] =
  [(SpritePropertyAtoms)]
  [(SpriteImageAtoms)]
```

```
[(SpriteOverrideSample)] =
  [(SpritePropertyAtoms)]
```

```
[(SpriteImageAtoms)]
  kSpriteSharedDataAtomType, 1, 1
    <kSpriteVariablesContainerAtomType>, 1
      <kSpriteStringVariableAtomType>, (1..n) ID is SpriteTrack
        Variable ID to be set
          [CString]
      <kSpriteFloatingPointVariableAtomType>, (1..n) ID is
```

Media Data Atom Types

```

        SpriteTrack Variable ID to be set
        [float]

kSpriteImagesContainerAtomType, 1, 1
    kSpriteImageAtomType, theImageID, (1 .. numImages)
        kSpriteImageDataAtomType, 1, 1
            [ImageData is ImageDescriptionHandle prepended to
                image data]

        <kSpriteImageRegistrationAtomType, 1, 1>
            [FixedPoint]
        <kSpriteImageNameAtomType, 1, 1>
            [pString]
        <kSpriteImageGroupIDAtomType, 1, 1>
            [long]

[(SpritePropertyAtoms)]
    <kQTEventFrameLoaded>, 1, 1
        [(ActionListAtoms)]
        <kCommentAtomType>, (anyUniqueIDs), (1..numComments)
            [CString]

kSpriteAtomType, theSpriteID, (1 .. numSprites)
    <kSpritePropertyMatrix, 1, 1>
        [MatrixRecord]
    <kSpritePropertyVisible, 1, 1>
        [short]
    <kSpritePropertyLayer, 1, 1>
        [short]
    <kSpritePropertyImageIndex, 1, 1>
        [short]
    <kSpritePropertyGraphicsMode, 1, 1>
        [ModifierTrackGraphicsModeRecord]

    <kSpriteUsesImageIDsAtomType, 1, 1>
        [array of QTAtomID's, one per image used]

    <kSpriteBehaviorsAtomType>, 1

    <kSpriteImageBehaviorAtomType>
        [QTSpriteButtonBehaviorStruct]

```

Media Data Atom Types

```

    <kSpriteCursorBehaviorAtomType>
        [QTSpriteButtonBehaviorStruct]
    <kSpriteStatusStringsBehaviorAtomType>
        [QTSpriteButtonBehaviorStruct]

    <[(SpriteActionAtoms)]>

[(SpriteActionAtoms)] =
    kQTEventType, theQTEventType, (1 .. numEventTypes)
    [(ActionListAtoms)] //see the next section Wired Action
                        //Grammar for a description
    <kCommentAtomType>, (anyUniqueIDs), (1..numComments)
    [CString]

```

Wired Action Grammar

The wired action grammar shown in this section allows QT event handlers to be expressed in a QuickTime movie. The sprite, text, VR, 3D, and Flash media handlers all support the embedding of QT event handlers in their media samples.

```

[(ActionListAtoms)] =
    kAction, (anyUniqueIDs), (1..numActions)
    kWhichAction    1, 1
    [long whichActionConstant]
    <kActionParameter> (anyUniqueIDs), (1..numParameters)
    [(parameterData)] ( whichActionConstant, paramIndex )
    // either leaf data or child atoms
    <kActionFlags> parameterID, (1..numParamsWithFlags)
    [long actionFlags]
    <kActionParameterMinValue> parameterID, (1.. numParamsWithMin)
    [data depends on param type]
    <kActionParameterMaxValue> parameterID, (1.. numParamsWithMax)
    [data depends on param type]
    [(ActionTargetAtoms)]

    <kCommentAtomType>, (anyUniqueIDs), (1..numComments)
    [CString]

```

Media Data Atom Types

```

[ (ActionTargetAtoms) ] =
    <kActionTarget>
        <kTargetMovie>
            [no data]
[ (ActionTargetAtoms) ] =
    <kActionTarget>

        <kTargetMovieName>
            [Pstring MovieName]
        OR
        <kTargetMovieID>
            [long MovieID]
        OR
        [ (kExpressionAtoms) ]
    <kTargetChildMovieTrackName>
        [PString childMovieTrackName]
    <kTargetChildMovieTrackID>
        [long childMovieTrackID]
    <kTargetChildMovieTrackIndex>
        [long childMovieTrackIndex]
    <kTargetChildMovieMovieName>
        [PString childMovieName]
    <kTargetChildMovieMovieID>
        [long childMovieID]
    <kTargetTrackName>
        [PString trackName]
    <kTargetTrackType>
        [OSType trackType]
    <kTargetTrackIndex>
        [long trackIndex]
        OR
        [ (kExpressionAtoms) ]
    <kTargetTrackID>
        [long trackID]
        OR
        [ (kExpressionAtoms) ]
    <kTargetSpriteName>
        [PString spriteName]
    <kTargetSpriteIndex>
        [short spriteIndex]
        OR

```


Media Data Atom Types

```

        [(kExpressionAtoms)]
        <kTargetSpriteID>
        [QTAtomID spriteIID]
        OR
        [(kExpressionAtoms)]
        <kTargetQD3DNamedObjectName>
        [CString objectName]

[(kExpressionAtoms)] =
    kExpressionContainerAtomType, 1, 1
    <kOperatorAtomType, theOperatorType, 1>
    kOperandAtomType, (anyUniqueIDs), (1..numOperands)
    [(OperandAtoms)]
    OR
    <kOperandAtomType, 1, 1>
    [(OperandAtoms)]

[(OperandAtoms)] =
    <kOperandExpression> 1, 1
    [(kExpressionAtoms)]           // allows for recursion
    OR
    <kOperandConstant> 1, 1
    [ float theConstant ]
    OR
    <kOperandSpriteTrackVariable> 1, 1
    [(ActionTargetAtoms)]
    kActionParameter, 1, 1
    [QTAtomID spriteVariableID]
    OR
    <kOperandKeyIsDown> 1, 1
    kActionParameter, 1, 1
    [UInt16 modifierKeys]
    kActionParameter, 2, 2
    [UInt8 asciiCharCode]
    OR
    <kOperandRandom> 1, 1
    kActionParameter, 1, 1
    [short minimum]
    kActionParameter, 2, 2
    [short maximum]

```

Media Data Atom Types

```
OR
<any other operand atom type>
  [(ActionTargetAtoms)]
```

The format for parameter data depends on the action and parameter index.

In most cases, the `kActionParameter` atom is a leaf atom containing data; for a few parameters, it contains child atoms.

whichAction corresponds to the action type that is specified by the leaf data of a `kWhichAction` atom.

paramIndex is the index of the parameter's `kActionParameter` atom.

```
[(parameterData)] ( whichAction, paramIndex ) =
{
    kActionMovieSetVolume:
        param1:      short volume

    kActionMovieSetRate
        param1:      Fixed rate

    kActionMovieSetLoopingFlags
        param1:      long loopingFlags

    kActionMovieGoToTime
        param1:      TimeValue time

    kActionMovieGoToTimeByName
        param1:      Str255 timeName

    kActionMovieGoToBeginning
        no params

    kActionMovieGoToEnd
        no params

    kActionMovieStepForward
        no params

    kActionMovieStepBackward
        no params
```

Media Data Atom Types

```
kActionMovieSetSelection
    param1:    TimeValue startTime
    param2:    TimeValue endTime

kActionMovieSetSelectionByName
    param1:    Str255 startTimeName
    param2:    Str255 endTimeName

kActionMoviePlaySelection
    param1:    Boolean selectionOnly

kActionMovieSetLanguage
    param1:    long language

kActionMovieChanged
    no params

kActionTrackSetVolume
    param1:    short volume

kActionTrackSetBalance
    param1:    short balance

kActionTrackSetEnabled
    param1:    Boolean enabled

kActionTrackSetMatrix
    param1:    MatrixRecord matrix

kActionTrackSetLayer
    param1:    short layer

kActionTrackSetClip
    param1:    RgnHandle clip

kActionSpriteSetMatrix
    param1:    MatrixRecord matrix

kActionSpriteSetImageIndex
    parml:    short imageIndex
```

Media Data Atom Types

```
kActionSpriteSetVisible
    param1:    short visible

kActionSpriteSetLayer
    param1:    short layer

kActionSpriteSetGraphicsMode
    param1:    ModifierTrackGraphicsModeRecord graphicsMode

kActionSpritePassMouseToCodec
    no params

kActionSpriteClickOnCodec
    param1:    Point localLoc

kActionSpriteTranslate
    param1:    Fixed x
    param2:    Fixed y
    param3:    Boolean isRelative

kActionSpriteScale
    param1:    Fixed xScale
    param2:    Fixed yScale

kActionSpriteRotate
    param1:    Fixed degrees

kActionSpriteStretch
    param1:    Fixed p1x
    param2:    Fixed p1y
    param3:    Fixed p2x
    param4:    Fixed p2y
    param5:    Fixed p3x
    param6:    Fixed p3y
    param7:    Fixed p4x
    param8:    Fixed p4y

kActionQTVRSetPanAngle
    param1:    float panAngle
```

Media Data Atom Types

```

kActionQTVRSetTiltAngle
    param1:    float tileAngle

kActionQTVRSetFieldOfView
    param1:    float fieldOfView

kActionQTVRShowDefaultView
    no params

kActionQTVRGoToNodeID
    param1:    UInt32 nodeID

kActionMusicPlayNote
    param1:    long sampleDescIndex
    param2:    long partNumber
    param3:    long delay
    param4:    long pitch
    param5:    long velocity
    param6:    long duration

kActionMusicSetController
    param1:    long sampleDescIndex
    param2:    long partNumber
    param3:    long delay
    param4:    long controller
    param5:    long value

kActionCase
    param1:    [(CaseStatementActionAtoms)]

kActionWhile
    param1:    [(WhileStatementActionAtoms)]

kActionGoToURL
    param1:    CString urlLink

kActionSendQTEventToSprite
    param1:    [(SpriteTargetAtoms)]
    param2:    QTEventRecord theEvent

kActionDebugStr

```

Media Data Atom Types

```

        param1:      Str255 theMessageString

kActionPushCurrentTime
    no params

kActionPushCurrentTimeWithLabel
    param1:      Str255 theLabel

kActionPopAndGotoTopTime
    no params

kActionPopAndGotoLabeledTime
    param1:      Str255 theLabel

kActionSpriteTrackSetVariable
    param1:      QTAtomID variableID
    param2:      float value

kActionApplicationNumberAndString
    param1:      long aNumber
    param2:      Str255 aString
}

```

Both `[(CaseStatementActionAtoms)]` and `[(WhileStatementActionAtoms)]` are child atoms of a `kActionParameter 1, 1` atom.

```

[(CaseStatementActionAtoms)] =
    kConditionalAtomType, (anyUniqueIDs), (1..numCases)
    [(kExpressionAtoms)]
    kActionListAtomType 1, 1
    [(ActionListAtoms)] // may contain nested conditional actions

[(WhileStatementActionAtoms)] =
    kConditionalAtomType, 1, 1
    [(kExpressionAtoms)]
    kActionListAtomType 1, 1
    [(ActionListAtoms)] // may contain nested conditional actions

```

Flash Media

Flash is a vector-based graphics and animation technology designed for the Internet. As an authoring tool, Flash lets content authors and developers create a wide range of interactive vector animations. The files exported by this tool are called SWF (pronounced “swiff”) files. SWF files are commonly played back using Macromedia’s ShockWave plug-in. In an effort to establish Flash as an industry-wide standard, Macromedia has published the SWF File Format and made the specification publicly available on its website at <http://www.macromedia.com/software/flash/open/spec/>.

The Flash media handler, introduced in QuickTime 4, allows a Macromedia Flash SWF 3.0 file to be treated as a track within a QuickTime movie. Thus, QuickTime 4 extends the SWF file format, enabling the execution of any of its wired actions. See “Adding Wired Actions To a Flash Track” (page 284) for an example of how to add wired actions.

Because a QuickTime movie may contain any number of tracks, multiple SWF tracks may be added to the same movie. The Flash media handler also provides support for an optimized case using the alpha channel graphics mode, which allows a Flash track to be composited cleanly over other tracks.

QuickTime supports all Flash actions except for the Flash load movie action. For example, when a Flash track in a QuickTime movie contains an action that goes to a particular Flash frame, QuickTime converts this to a wired action that goes to the QuickTime movie time in the corresponding Flash frame.

Note

As a time-based media playback format, QuickTime may drop frames when necessary to maintain its schedule. As a consequence, frames of a SWF file may be dropped during playback. If this is not satisfactory for your application, you can set the playback mode of the movie to Play All Frames, which will emulate the playback mode of ShockWave. QuickTime’s SWF file importer sets the Play All Frames mode automatically when adding a SWF file to an empty movie. ♦

QuickTime support for Flash 3.0 also includes the DoFSCommand mechanism. This allows JavaScript routines with a specific function prototype to be invoked with

parameters passed from the Flash track. Refer to Macromedia's Flash 3 documentation for more details on how to author a .SWF 3.0 file with a Flash FSCCommand.

Tween Media

Tween media is used to store pairs of values to be interpolated between in QuickTime movies. These interpolated values modify the playback of other media types by using track references and track input maps. For example, a tween media could generate gradually changing relative volume levels to cause an audio track to fade out. It has a media type of 'tween'.

Every **tween** operation is based on a collection of one or more values from which a range of output values can be algorithmically derived. Each tween is assigned a time duration, and an output value can be generated for any time value within the duration. In the simplest kind of tween operation, a pair of values is provided as input and values *between* the two values are generated as output.

A **tween track** is a special track in a movie that is used exclusively as a modifier track. The data it contains, known as **tween data**, is used to generate values that modify the playback of other tracks, usually by interpolating values. The tween media handler sends these values to other media handlers; it never presents data.

Tween Sample Description

The tween sample description uses the standard sample description header, as described in “Sample Table Atoms” (page 85).

The data format field in the sample description is always set to 'tween'. The tween media handler adds no additional fields to the sample description.

Tween Sample Data

Tween sample data is stored in QT atom structures.

At the root level, there are one or more tween entry atoms; these atoms have an atom type value of 'tween'. Each tween entry atom completely describes one

Media Data Atom Types

interpolation operation. These atoms should be consecutively numbered starting at 1, using the atom ID field.

Each tween entry atom contains several more atoms that describe how to perform the interpolation. The atom ID field in each of these atoms must be set to 1.

- Tween start atom (atom type is 'twst'). This atom specifies the time at which the interpolation is to start. The time is expressed in the media's time coordinate system. If this atom is not present, the starting offset is assumed to be 0.
- Tween duration atom (atom type is 'twdu'). This atom specifies how long the interpolation is to last. The time is expressed in the media's time coordinate system. If this atom is not present, the duration is assumed to be the length of the sample.
- Tween data atom (atom type is 'twdt'). This atom contains the actual values for the interpolation. The contents depend on the value of the tween type atom.
- Tween type atom (atom type is 'twnt'). Describes the type of interpolation to perform. Table 3-7 shows all currently defined tween types. All tween types are currently supported using linear interpolation.

Table 3-7 Tween type values

Tween type	Value	Tween data
16-bit integer	1	Two 16-bit integers.
32-bit integer	2	Two 32-bit integers.
32-bit fixed-point	3	Two 32-bit fixed-point numbers.
Point: two 16-bit integers	4	Two points.
Rectangle: four 16-bit integers	5	Two rectangles.

Table 3-7 Tween type values

Tween type	Value	Tween data
QuickDraw region	6	Two rects and a region. The tween entry atom must contain a 'qdrg' atom with an atom ID value of 1. The region is transformed through the resulting matrices.
Matrix	7	Two matrices.
RGB color: three 16-bit integers	8	Two RGB colors.
Graphics mode with RGB color	9	Two graphics modes with RGB color. Only the RGB color is interpolated. The graphics modes must be the same.

Each tween type is distinguished from other types by these characteristics:

- input values or structures of a particular type
- a particular number of input values or structures (most often one or two)
- output values or structures of a particular type
- a particular algorithm used to derive the output values

Tween operations for each tween type are performed by a tween component that is specific to that type or, for a number of tween types that are native to QuickTime, by QuickTime itself. Movies and applications that use tweening do not need to specify the tween component to use; QuickTime identifies a tween type by its tween type identifier and automatically routes its data to the correct tween component or to QuickTime.

When a movie contains a tween track, the tween media handler invokes the necessary component (or built-in QuickTime code) for tween operations and delivers the results to another media handler. The receiving media handler can then use the values it receives to modify its playback. For example, the data in a tween track can be used to alter the volume of a sound track.

Tweening can also be used outside of movies by applications or other software that can use the values it generates.

Tween Type Categories

Each of the tween types supported by QuickTime belongs to one of these categories:

- Numeric tween types, which have pairs of numeric values, such as long integers, as input. For these types, linear interpolation is used to generate output values.
- QuickDraw tween types, most of which have pairs of QuickDraw structures, such as points or rectangle, as input. For these types, one or more structure elements are interpolated, such as the *h* and *v* values for points, and each element that is interpolated is interpolated separately from others.
- 3D tween types, which have a QuickDraw 3D structure such as *TQ3Matrix4x4* or *TQ3RotateAboutAxisTransformData* as input. For these types, a specific 3D transformation is performed on the data to generate output.
- The polygon tween type, which takes three four-sided polygons as input. One polygon (such as the bounds for a sprite or track) is transformed, and the two others specify the start and end of the range of polygons into which the tween operation maps it. You can use the output (a *MatrixRecord* data structure) to map the source polygon into any intermediate polygon. The intermediate polygon is interpolated from the start and end polygons for each particular time in the tween duration.
- Path tween types have as input a QuickTime vector data stream for a path. Four of the path tween types also have as input a percentage of path's length; for these types, either a point on the path or a data structure is returned. Two other path tween types treat the path as a function: one returns the *y* value of the point on the path with a given *x* value, and the other returns the *x* value of the point on the path with a given *y* value.
- The list tween type has as input a QT atom container that contains leaf atoms of a specified atom type. For this tween type category, the duration of the tween operation is divided by the number of leaf atoms of the specified type. For time points within the first time division, the data for the first leaf atom is returned; for the second time division, the data for the second leaf atom is returned; and so on. The resulting tween operation proceeds in discrete steps (one step for each leaf atom), instead of the relatively continuous tweening produced by other tween type categories.

Tween QT Atom Container

The characteristics of a tween are specified by the atoms in a tween QT atom container.

A tween QT atom container can contain the atoms described in the following sections.

General Tween Atoms

■ `kTweenEntry`

Specifies a tween atom, which can be either a single tween atom, a tween atom in a tween sequence, or an interpolation tween atom.

Its parent is the tween QT atom container (which you specify with the constant `kParentAtomIsContainer`).

The index of a `kTweenEntry` atom specifies when it was added to the QT atom container; the first added has the index 1, the second 2, and so on. The ID of a `kTweenEntry` atom can be any ID that is unique among the `kTweenEntry` atoms contained in the same QuickTime atom container.

This atom is a parent atom. It must contain the following child atoms:

- A `kTweenType` atom that specifies the tween type.
- One or more `kTweenData` atoms that contain the data for the tween atom. Each `kTweenData` atom can contain different data to be processed by the tween component, and a tween component can process data from only one `kTweenData` atom a time. For example, an application can use a list tween to animate sprites. The `kTweenEntry` atom for the tween atom could contain three sets of animation data, one for moving the sprite from left to right, one for moving the sprite from right to left, and one for moving the sprite from top to bottom. In this case, the `kTweenEntry` atom for the tween atom would contain three `kTweenData` atoms, one for each data set. The application specifies the desired data set by specifying the ID of the `kTweenData` atom to use.

A `kTweenEntry` atom can contain any of the following optional child atoms:

- A `kTweenStartOffset` atom that specifies a time interval, beginning at the start of the tween media sample, after which the tween operation begins. If this atom is not included, the tween operation begins at the start of the tween media sample.

Media Data Atom Types

- A `kTweenDuration` atom that specifies the duration of the tween operation. If this atom is not included, the duration of the tween operation is the duration of the media sample that contains it.

If a `kTweenEntry` atom specifies a path tween, it can contain the following optional child atom:

- A `kTweenFlags` atom containing flags that control the tween operation. If this atom is not included, no flags are set.

Note that **interpolation tween tracks** are tween tracks that modify other tween tracks. The output of an interpolation tween track must be a time value, and the time values generated are used in place of the input time values of the tween track being modified.

If a `kTweenEntry` atom specifies an interpolation tween track, it must contain the following child atoms:

- A `kTweenInterpolationID` atom for each `kTweenData` atom to be interpolated. The ID of each `kTweenInterpolationID` atom must match the ID of the `kTweenData` atom to be interpolated. The data for a `kTweenInterpolationID` atom specifies a `kTweenEntry` atom that contains the interpolation tween track to use for the `kTweenData` atom.

If this atom specifies an interpolation tween track, it can contain either of the following optional child atoms:

- A `kTweenOutputMin` atom that specifies the minimum output value of the interpolation tween atom. The value of this atom is used only if there is also a `kTweenOutputMax` atom with the same parent. If this atom is not included and there is a `kTweenOutputMax` atom with the same parent, the tween component uses 0 as the minimum value when scaling output values of the interpolation tween track.
- A `kTweenOutputMax` atom that specifies the maximum output value of the interpolation tween atom. If this atom is not included, the tween component does not scale the output values of the interpolation tween track.

■ `kTweenStartOffset`

For a tween atom in a tween track of a QuickTime movie, specifies a time offset from the start of the tween media sample to the start of the tween atom. The time units are the units used for the tween track.

Its parent atom is a `kTweenEntry` atom.

A `kTweenEntry` atom can contain only one `kTweenStartOffset` atom. The ID of this atom is always 1. The index of this atom is always 1.

This atom is a leaf atom. The data type of its data is `TimeValue`.

Media Data Atom Types

This atom is optional. If it is not included, the tween operation begins at the start of the tween media sample.

■ `kTweenDuration`

Specifies the duration of a tween operation. When a QuickTime movie includes a tween track, the time units for the duration are those of the tween track. If a tween component is used outside of a movie, the application using the tween data determines how the duration value and values returned by the component are interpreted.

Its parent atom is a `kTweenEntry` atom.

A `kTweenEntry` atom can contain only one `kTweenDuration` atom. The ID of this atom is always 1. The index of this atom is always 1.

This atom is a leaf atom. The data type of its data is `TimeValue`.

This atom is optional. If it is not included, the duration of the tween operation is the duration of the media sample that contains it.

■ `kTweenData`

Contains data for a tween atom.

Its parent atom is a `kTweenEntry` atom.

A `kTweenEntry` atom can contain any number of `kTweenData` atoms.

The index of a `kTweenData` atom specifies when it was added to the `kTweenEntry` atom; the first added has the index 1, the second 2, and so on.

The ID of a `kTweenData` atom can be any ID that is unique among the `kTweenData` atoms contained in the same `kTweenEntry` atom.

At least one `kTweenData` atom is required in a `kTweenEntry` atom.

For single tween atoms, a `kTweenData` atom is a leaf atom. It can contain data of any type.

For polygon tween atoms, a `kTweenData` atom is a leaf atom. The data type of its data is `Fixed[27]`, which specifies three polygons.

For path tweens, a `kTweenData` atom is a leaf atom. The data type of its data is `Handle`, which contains a QuickTime vector.

In interpolation tween atoms, a `kTweenData` atom is a leaf atom. It can contain data of any type. An interpolation tween atom can be any tween atoms other than a list tween atom that returns a time value.

In list tween atoms, a `kTweenData` atom is a parent atom that must contain the following child atoms:

- A `kListElementType` atom that specifies the atom type of the elements of the tween atom.

Media Data Atom Types

- One or more leaf atoms of the type specified by the `kListElementType` atom. The data for each atom is the result of a list tween operation.
- `kNameAtom`

Specifies the name of a tween atom. The name, which is optional, is not used by tween components, but it can be used by applications or other software.

Its parent atom is a `kTweenEntry` atom.

A `kTweenEntry` atom can contain only one `kNameAtom` atom. The ID of this atom is always 1. The index of this atom is always 1.

This atom is a leaf atom. Its data type is `String`.

This atom is optional. If it is not included, the tween atom does not have a name.
- `kTweenType`

Specifies the tween type (the data type of the data for the tween operation).

Its parent atom is a `kTweenEntry` atom.

A `kTweenEntry` atom can contain only one `kTweenType` atom. The ID of this atom is always 1. The index of this atom is always 1.

This atom is a leaf atom. The data type of its data is `OSType`.

This atom is required.

Path Tween Atoms

- `kTweenFlags`

Contains flags that control the tween operation. One flag that controls path tween atoms is defined:

 - The `kTweenReturnDelta` flag applies only to path tween atoms (tweens of type `kTweenTypePathToFixedPoint`, `kTweenTypePathToMatrixTranslation`, `kTweenTypePathToMatrixTranslationAndRotation`, `kTweenTypePathXtoY`, or `kTweenTypePathYtoX`). If the flag is set, the tween component returns the change in value from the last time it was invoked. If the flag is not set, or if the tween component has not previously been invoked, the tween component returns the normal result for the tween atom.

Its parent atom is a `kTweenEntry` atom.

A `kTweenEntry` atom can contain only one `kTweenFlags` atom. The ID of this atom is always 1. The index of this atom is always 1.

This atom is a leaf atom. The data type of its data is `Long`.

This atom is optional. If it is not included, no flags are set.

Media Data Atom Types

■ `kInitialRotationAtom`

Specifies an initial angle of rotation for a path tween atom of type `kTweenTypePathToMatrixRotation`, `kTweenTypePathToMatrixTranslation`, or `kTweenTypePathToMatrixTranslationAndRotation`.

Its parent atom is a `kTweenEntry` atom.

A `kTweenEntry` atom can contain only one `kInitialRotationAtom` atom. The ID of this atom is always 1. The index of this atom is always 1.

This atom is a leaf atom. Its data type is `Fixed`.

This atom is optional. If it is not included, no initial rotation of the tween atom is performed.

List Tween Atoms

■ `kListElementType`

Specifies the atom type of the elements in a list tween atom.

Its parent atom is a `kTweenData` atom.

A `kTweenEntry` atom can contain only one `kListElementType` atom. The ID of this atom is always 1. The index of this atom is always 1.

This atom is a leaf atom. Its data type is `QTAtomType`.

This atom is required in the `kTweenData` atom for a list tween atom.

3D Tween Atoms

■ `kTween3dInitialCondition`

Specifies an initial transform for a 3D tween atom whose tween type is one of the following: `kTweenType3dCameraData`, `kTweenType3dMatrix`, `kTweenType3dQuaternion`, `kTweenType3dRotate`, `kTweenType3dRotateAboutAxis`, `kTweenType3dRotateAboutAxis`, `kTweenType3dRotateAboutPoint`, `kTweenType3dRotateAboutVector`, `kTweenType3dScale`, or `kTweenType3dTranslate`.

Its parent atom is a `kTweenEntry` atom.

A `kTweenEntry` atom can contain only one `kTween3dInitialCondition` atom. The ID of this atom is always 1. The index of this atom is always 1.

This atom is a leaf atom. The data type of its data is as follows:

- ☐ For a `kTweenType3dCameraData` tween, its data type is `TQ3CameraData`.
- ☐ For a `kTweenType3dMatrix` tween, its data type is `TQ3Matrix4x4`.

Media Data Atom Types

- For a `kTweenType3dQuaternion` tween, its data type is `TQ3Quaternion`.
- For a `kTweenType3dRotate` tween, its data type is `TQ3RotateTransformData`.
- For a `kTweenType3dRotateAboutAxis` tween, its data type is `TQ3RotateAboutAxisTransformData`.
- For a `kTweenType3dRotateAboutPoint` tween, its data type is `TQ3RotateAboutPointTransformData`.
- For a `kTweenType3dRotateAboutVector` tween, its data type is `TQ3PlaneEquation`.
- For a `kTweenType3dScale` tween, its data type is `TQ3Vector3D`.
- For a `kTweenType3dTranslate` tween, its data type is `TQ3Vector3D`.

This atom is optional. For each tween type, the default value is the data structure that specifies an identity transform, that is, a transform that does not alter the 3D data.

Interpolation Tween Atoms

■ `kTweenOutputMax`

Specifies the maximum output value of an interpolation tween atom. If a `kTweenOutputMax` atom is included for an interpolation tween, output values for the tween atom are scaled to be within the minimum and maximum values. The minimum value is either the value of the `kTweenOutputMin` atom or, if there is no `kTweenOutputMin` atom, 0. For example, if an interpolation tween atom has values between 0 and 4, and it has `kTweenOutputMin` and `kTweenOutputMax` atoms with values 1 and 2, respectively, a value of 0 (the minimum value before scaling) is scaled to 1 (the minimum specified by the `kTweenOutputMin` atom), a value of 4 (the maximum value before scaling) is scaled to 2 (the maximum specified by the `kTweenOutputMax` atom), and a value of 3 (three-quarters of the way between the maximum and minimum values before scaling) is scaled to 1.75 (three-quarters of the way between the values of the `kTweenOutputMin` and `kTweenOutputMax` atoms).

Its parent atom is a `kTweenEntry` atom.

A `kTweenEntry` atom can contain only one `kTweenOutputMax` atom. The ID of this atom is always 1. The index of this atom is always 1.

This atom is a leaf atom. The data type of its data is `Fixed`.

This atom is optional. If it is not included, QuickTime does not scale interpolation tween values.

■ `kTweenOutputMin`

Specifies the minimum output value of an interpolation tween atom. If both `kTweenOutputMin` and `kTweenOutputMax` atoms are included for an interpolation tween atom, output values for the tween atom are scaled to be within the minimum and maximum values. For example, if an interpolation tween atom has values between 0 and 4, and it has `kTweenOutputMin` and `kTweenOutputMax` atoms with values 1 and 2, respectively, a value of 0 (the minimum value before scaling) is scaled to 1 (the minimum specified by the `kTweenOutputMin` atom), a value of 4 (the maximum value before scaling) is scaled to 2 (the maximum specified by the `kTweenOutputMax` atom), and a value of 3 (three-quarters of the way between the maximum and minimum values before scaling) is scaled to 1.75 (three-quarters of the way between the values of the `kTweenOutputMin` and `kTweenOutputMax` atoms).

If a `kTweenOutputMin` atom is included but a `kTweenOutputMax` atom is not, QuickTime does not scale interpolation tween values.

Its parent atom is a `kTweenEntry` atom.

A `kTweenEntry` atom can contain only one `kTweenOutputMin` atom. The ID of this atom is always 1. The index of this atom is always 1.

This atom is a leaf atom. The data type of its data is `Fixed`.

This atom is optional. If it is not included but a `kTweenOutputMax` atom is, the tween component uses 0 as the minimum value for scaling values of an interpolation tween atom.

■ `kTweenInterpolationID`

Specifies an interpolation tween atom to use for a specified `kTweenData` atom. There can be any number of `kTweenInterpolationID` atoms for a tween atom, one for each `kTweenData` atom to be interpolated.

Its parent atom is a `kTweenEntry` atom.

The index of a `kTweenInterpolationID` atom specifies when it was added to the `kTweenEntry` atom; the first added has the index 1, the second 2, and so on. The ID of a `kTweenInterpolationID` atom must match the atom ID of the `kTweenData` atom to be interpolated, and be unique among the `kTweenInterpolationID` atoms contained in the same `kTweenEntry` atom.

This atom is a leaf atom. The data type of its data is `QTAtomID`.

This atom is required for an interpolation tween atom.

Region Tween Atoms

■ `kTweenPictureData`

Media Data Atom Types

Contains the data for a QuickDraw picture. Used only by a `kTweenTypeQDRegion` atom.

Its parent atom is a `kTweenEntry` atom.

A `kTweenEntry` atom can contain only one `kTweenPictureData` or `kTweenRegionData` atom. The ID of this atom is always 1. The index of this atom is always 1.

This atom is a leaf atom. The data type of its data is `Picture`.

Either a `kTweenPictureData` or `kTweenRegionData` atom is required for a `kTweenTypeQDRegion` atom.

■ `kTweenRegionData`

Contains the data for a QuickDraw region. Used only by a `kTweenTypeQDRegion` atom.

Its parent atom is a `kTweenEntry` atom.

A `kTweenEntry` atom can contain only one `kTweenRegionData` or `kTweenPictureData` atom. The ID of this atom is always 1. The index of this atom is always 1.

This atom is a leaf atom. The data type of its data is `Region`.

Either a `kTweenPictureData` or `kTweenRegionData` atom is required for a `kTweenTypeQDRegion` tween.

Sequence Tween Atoms

■ `kTweenSequenceElement`

Specifies an entry in a tween sequence.

Its parent is the tween QT atom container (which you specify with the constant `kParentAtomIsContainer`).

The ID of a `kTweenSequenceElement` atom must be unique among the `kTweenSequenceElement` atoms in the same QT atom container. The index of a `kTweenSequenceElement` atom specifies its order in the sequence; the first entry in the sequence has the index 1, the second 2, and so on.

This atom is a leaf atom. The data type of its data is

`TweenSequenceEntryRecord`, a data structure that contains the following fields:

<code>endPercent</code>	A value of type <code>Fixed</code> that specifies the point in the duration of the tween media sample at which the sequence entry ends. This is expressed as a percentage; for example, if the value is 75.0, the sequence entry ends after three-quarters of the
-------------------------	---

total duration of the tween media sample have elapsed. The sequence entry begins after the end of the previous sequence entry or, for the first entry in the sequence, at the beginning of the tween media sample.

<code>tweenAtomID</code>	A value of type <code>QTAtomID</code> that specifies the <code>kTweenEntry</code> atom containing the tween for the sequence element. The <code>kTweenEntry</code> atom and the <code>kTweenSequenceElement</code> atom must both be a child atoms of the same tween QT atom container.
<code>dataAtomID</code>	A value of type <code>QTAtomID</code> that specifies the <code>kTweenData</code> atom containing the data for the tween. This atom must be a child atom of the atom specified by the <code>tweenAtomID</code> field.

Modifier Tracks

The addition of **modifier tracks** in QuickTime 2.1 introduced the capability for creating dynamic movies. (A modifier track sends data to another track; by comparison, a track reference is an association.) For example, instead of playing video in a normal way, a video track could send its image data to a sprite track. The sprite track then could use that video data to replace the image of one of its sprites. When the movie is played, the video track appears as a sprite.

Modifier tracks are not a new type of track. Instead, they are a new way of using the data in existing tracks. A modifier track does not present its data, but sends it to another track that uses the data to modify how it presents its own data. Any track can be either a sender or a presenter, but not both. Previously, all tracks were presenters.

Another use of modifier tracks is to store a series of sound volume levels, which is what occurs when you work with a tween track. These sound levels can be sent to a sound track as it plays to dynamically adjust the volume. A similar use of modifier tracks is to store location and size information. This data can be sent to a video track to cause it to move and resize as it plays.

Because a modifier track can send its data to more than one track, you can easily synchronize actions between multiple tracks. For example, a single

modifier track containing matrices as its samples can make two separate video tracks follow the same path.

See “Creating Movies With Modifier Tracks” (page 280) for more information about using modifier tracks.

Limitations of Spatial Modifier Tracks

A modifier track may cause a track to move outside of its original boundary regions. This may present problems, since applications do not expect the dimensions or location of a QuickTime movie to change over time.

To ensure that a movie maintains a constant location and size, the Movie Toolbox limits the area in which a spatially modified track can be displayed. A movie’s “natural” shape is defined by the region returned by the `GetMovieBoundsRgn` function. The toolbox clips all spatially modified tracks against the region returned by `GetMovieBoundsRgn`. This means that a track can move outside of its initial boundary regions, but it cannot move beyond the combined initial boundary regions of all tracks in the movie. Areas uncovered by a moving track are handled by the toolbox in the same way as areas uncovered by tracks with empty edits.

If a track has to move through a larger area than that defined by the movie’s boundary region, the movie’s boundary region can be enlarged to any desired size by creating a spatial track (such as a video track) of the desired size but with no data. As long as the track is enabled, it contributes to the boundary regions of the movie.

Track References

Although QuickTime has always allowed the creation of movies that contain more than one track, it has not been able to specify relationships between those tracks. **Track references** are a feature of QuickTime that allows you to relate a movie’s tracks to one another. The QuickTime track-reference mechanism supports many-to-many relationships. That is, any movie track may contain one or more track references, and any track may be related to one or more other tracks in the movie.

Track references can be useful in a variety of ways. For example, track references can be used to relate timecode tracks to other movie tracks. You can

use track references to identify relationships between video and sound tracks—identifying the track that contains dialog and the track that contains background sounds, for example. Another use of track references is to associate one or more text tracks that contain subtitles with the appropriate audio track or tracks.

Track references are also used to create chapter lists, as described in the next section.

Every movie track contains a list of its track references. Each track reference identifies another related track. That related track is identified by its track identifier. The track reference itself contains information that allows you to classify the references by type. This type information is stored in an `OSType` data type. You are free to specify any type value you want. Note, however, that Apple has reserved all lowercase type values.

You may create as many track references as you want, and you may create more than one reference of a given type. Each track reference of a given type is assigned an index value. The index values start at 1 for each different reference type. The Movie Toolbox maintains these index values, so that they always start at 1 and count by 1.

Using the `AddTrackReference` function, you can relate one track to another. The `DeleteTrackReference` function will remove that relationship. The `SetTrackReference` and `GetTrackReference` functions allow you to modify an existing track reference so that it identifies a different track. The `GetNextTrackReferenceType` and `GetTrackReferenceCount` functions allow you to scan all of a track's track references.

Chapter Lists

A **chapter list** provides a set of named entry points into a movie, allowing the user to jump to a preselected point in the movie from a convenient pop-up list.

The movie controller automatically recognizes a chapter list and will create a pop-up list from it. When the user makes a selection from the pop-up, the controller will jump to the appropriate point in the movie. Note that if the movie is sized so that the controller is too narrow to display the chapter names, the pop-up list will not appear.

To create a chapter list, you must create a text track with one sample for each chapter. The display time for each sample corresponds to the point in the movie that marks the beginning of that chapter. You must also create a track reference

Media Data Atom Types

of type 'chap' from an enabled track of the movie to the text track. It is the 'chap' track reference that makes the text track into a chapter list. The track containing the reference can be of any type (audio, video, MPEG, and so on), but it must be enabled for the chapter list to be recognized.

Given an enabled track `myVideoTrack`, for example, you can use the `AddTrackReference` function to create the chapter reference:

```
AddTrackReference( myVideoTrack, theTextTrack,
                   kTrackReferenceChapterList,
                   &addedIndex );
```

`kTrackReferenceChapterList` is defined in `Movies.h`. It has the value 'chap'.

The text track that constitutes the chapter list does not need to be enabled, and normally is not. If it is enabled, the text track will be displayed as part of the movie, just like any other text track, in addition to functioning as a chapter list.

If more than one enabled track includes a 'chap' track reference, QuickTime uses the first chapter list that it finds.

3D Media

QuickTime movies store 3D image data in a base media. This media has a media type of 'qd3d'.

3D Sample Description

The 3D sample description uses the standard sample description header, as described in “Sample Table Atoms” (page 85).

The data format field in the sample description is always set to 'qd3d'. The 3D media handler adds no additional fields to the sample description.

3D Sample Data

The 3D samples are stored in the 3D Metafile format developed for QuickDraw 3D.

Streaming Media

QuickTime movies store streaming data in a streaming media track. This media has a media type of 'strm'.

Streaming Media Sample Description

The streaming media sample description contains information that defines how to interpret streaming media data. This sample description is based on the standard sample description header, as described in “Sample Table Atoms” (page 85).

The streaming media sample description is documented in the QuickTime header file `QTSMovie.h`, as shown in Listing 3-1.

Listing 3-1 Streaming media sample description

```
struct QTSSampleDescription {
    long                descSize;
    long                dataFormat;
    long                resvd1;    /* set to 0*/
    short               resvd2;    /* set to 0*/
    short               dataRefIndex;
    UInt32              version;
    UInt32              resvd3;    /* set to 0*/
    SInt32              flags;

                                /* qt atoms follow*/
                                /* long size, long type, some data*/
                                /* repeat as necessary*/
};

typedef struct QTSSampleDescription    QTSSampleDescription;
```

The sample format depends on the `dataFormat` field of the `QTSSampleDescription`. The `dataFormat` field can be any value you specify. The currently defined values are 'rtsp' and 'sdp'.

Media Data Atom Types

If 'rtsp', the sample can be just an rtsp URL. It can also be any value that you can put in a .rtsp file, as defined at

<http://streaming.apple.com/qtstreaming/documentation/userdocs/rtsp-tags.htm>

If 'sdp', then the sample is an SDP file. This would be used to receive a multicast broadcast.

Hint Media

The QuickTime file format supports streaming of media data over a network as well as local playback. The process of sending protocol data units is time-based, just like the display of time-based data, and is therefore suitably described by a time-based format. A QuickTime file or movie that supports streaming includes information about the data units to stream. This information is included in additional tracks of the movie called **hint tracks**.

Hint tracks contain instructions for a streaming server which assist in the formation of packets. These instructions may contain immediate data for the server to send (for example, header information) or reference segments of the media data. These instructions are encoded in the QuickTime file in the same way that editing or presentation information is encoded in a QuickTime file for local playback.

Instead of editing or presentation information, information is provided which allows a server to packetize the media data in a manner suitable for streaming, using a specific network transport.

The same media data is used in a QuickTime file which contains hints, whether it is for local playback, or streaming over a number of different transport types. Separate hint tracks for different transport types may be included within the same file and the media will play over all such transport types without making any additional copies of the media itself. In addition, existing media can be easily made streamable by the addition of appropriate hint tracks for specific transports. The media data itself need not be recast or reformatted in any way.

Typically, hinting is performed by media packetizer components. QuickTime selects an appropriate media packetizer for each track and routes each packetizer's output through an Apple-provided packet builder to create a hint track. One hint track is created for each streamable track in the movie.

Hint tracks are quite small compared with audio or video tracks. A movie that contains hint tracks can be played from a local disk or streamed over HTTP, similar to any other QuickTime movie. Hint tracks are only used when streaming a movie over a real-time media streaming protocol, such as RTP.

Support for streaming in the QuickTime file format is based upon the following considerations:

- Media data represented as a set of network-independent standard QuickTime tracks, which may be played or edited, as normal.
- A common declaration and base structure for server hint tracks; this common format is protocol independent, but contains the declarations of which protocols are described in the server tracks.
- A specific design of the server hint tracks for each protocol which may be transmitted; all these designs use the same basic structure.

The resulting streams, sent by the servers under the direction of hint tracks, do not need to contain any trace of QuickTime information. This approach does not require that QuickTime, or its structures or declaration style, be used either in the data on the wire or in the decoding station. For example, a QuickTime file using H.261 video and DVI audio, streamed under Real-Time Protocol (RTP), results in a packet stream which is fully compliant with the IETF specifications for packing those codings into RTP.

Hint tracks are built and flagged, so that when the movie is viewed directly (not streamed), they are ignored.

The next section describes a generic format for streaming hints to be stored in a QuickTime movie.

Adding Hint Tracks to a Movie

To store packetization hints, one or more hint tracks are added to a movie. Each hint track contains hints for at least one actual media track to be streamed. A streamed media track may have more than one hint track. For example, it might have a separate hint track for the different packet sizes the server supports, or it might have different hint tracks for different protocols. It is not required that all media tracks have corresponding hint tracks in a movie.

The sample time of a hint sample corresponds to the sample time of the media contained in the packets generated by that hint sample. The hint sample may

also contain a transmission time for each packet. (The format for the hint sample is specific to the hint track type.)

The hint track may have a different time scale than its referenced media tracks.

The `flags` field in the track header atom ('`tkhd`') must be set to `0x000000`, indicating that the track is inactive and is not part of the movie, preview, or poster.

The `subType` field of the handler description atom ('`hdlr`') contains '`hint`', indicating that the media type is packetization hints.

Note that if a QuickTime media track is edited, any previously stored packetization hints may become invalid. Comparing the modification dates of the media track and the hint track is one way to determine this scenario, but is far from being foolproof. Since the hint track keeps track of which original track media samples and sample descriptions to play at specific times, changes that affect those parts of the original track or media make those hints invalid. Changes to a movie that do not invalidate existing hint tracks include flattening (when there are no edit lists), and adding new tracks. Changes that invalidate hint tracks include

- flattening (when there are edit lists)
- adding or deleting samples
- changing a track's time scale
- changing sample descriptions

▲ WARNING

Hint tracks are marked as inactive, so calling the `FlattenMovie` function with the `flattenActiveTracksOnly` bit set deletes all hint tracks from a movie. ▲

Packetization Hint Media Header Atom

In QuickTime movies, the media information atom ('`minf`') contains header data specific to the media. For hint tracks, the media header is a base media information atom ('`gmhd`'). The hint track must contain the base media information atom.

Hint Track User Data Atom

Each hint track may contain track user data atoms that apply to only to the corresponding hint track. There are currently two such atoms defined.

- User data atom type `'hinf'`.

This contains statistics for the hint track. The `'hinf'` atom contains child atoms as defined in Table 3-8. In some cases, there are both 32-bit and 64-bit counters available. Any unknown types should be ignored.

- User data atom type `'hnti'`.

This may contain child atoms. Child atoms that start with `'sdp '` (note, again, the space) contain SDP text for this track. Text from these child atoms must be inserted into the proper place in the SDP text for the movie, after any common SDP text. This is analogous to the movie-level `'hnti'` atom.

Movie Hint Info Atom

A movie may contain a `'hnti'` movie user data atom, which may contain one or more child atoms. The child atom contents start with 4 bytes that specify the transport and 4 bytes that specify the type of data contained in the rest of the child atom. Currently, the only defined transport is `'rtp '` (note the space) and the only content data type defined is `'sdp '` (note the space). Child atoms whose transport or type combinations you don't recognize should be skipped.

The text in an atom of type `'rtp sdp '` should be inserted (in the proper place) into the SDP information generated from this file (for example, by a streaming server) before any SDP information for specific tracks.

Table 3-8 The 'hinf' atom type containing child atoms

Type	Value	Description
'trpy'	8 bytes	The total number of bytes that will be sent, including 12-byte RTP headers, but not including any network headers.
'totl'	4 bytes	4-byte version of 'trpy'
'nump'	8 bytes	The total number of network packets that will be sent (if the application knows there is a 28-byte network header, it can multiply 28 by this number and add it to the 'trpy' value to get the true number of bytes sent.
'npck'	4 bytes	4-byte version of 'nump'
'tpyl'	8 bytes	The total number of bytes that will be sent, not including 12-byte RTP headers.
'tpay'	4 bytes	4-byte version of 'tpyl'
'maxr'	8 bytes	The maximum data rate. This atom contains two numbers: g, followed by m (both 32-bit values). g is the granularity, in milliseconds. m is the maximum data rate given that granularity. For example, if g is 1 second, then m is the maximum data rate over any 1 second. There may be multiple 'maxr' atoms, with different values for g. The maximum data rate calculation does not include any network headers (but does include 12-byte RTP headers).
'dmed'	8 bytes	The number of bytes from the media track to be sent.
'dimm'	8 bytes	The number of bytes of immediate data to be sent.
'drep'	8 bytes	The number of bytes of repeated data to be sent.
'tmin'	4 bytes	The smallest relative transmission time, in milliseconds.
'tmax'	4 bytes	The largest relative transmission time, in milliseconds.

Media Data Atom Types

Type	Value	Description
'pmax'	4 bytes	The largest packet, in bytes; includes 12-byte RTP header.
'dmax'	4 bytes	The largest packet duration, in milliseconds.
'payt'	variable	The payload type, which includes payload number (32-bits) followed by <code>rtptime</code> payload string (Pascal string).

Note

Any of the atoms shown in Table 3-8 may or may not be present. These atoms are not guaranteed. ♦

Finding an Original Media Track From a Hint Track

Like any other QuickTime track, hint tracks can contain track reference atoms. Exactly one of these must be of track reference type `'hint'`, and its internal list must contain at least one track ID, which is the track ID of the original media track. Like other track reference atoms, there may be empty references in this list, indicated by a track ID of 0. For hint tracks that refer to more than one track, the index number (starting at 1, and including any 0 entries) is used in the media track reference index field in some of the packet data table entry modes.

For example, if you have MPEG-1 video at track ID 11 and MPEG-1 layer 2 audio at track ID 12, and you are creating a RTP hint track that encapsulates these in an MPEG-2 transport, you need to refer to both tracks. You can also assume that there are some empty entries and other track references in your hint track atom reference atom's list. So it might look like this: 11, 0, 0, 14, 0, 12, 0. When you are assembling packets from audio and video tracks 11 and 12, you use their list indexes (1 and 6) in the media track ref index field.

If you have only one media track listed in your hint track reference, you may simply use a 0 in the media track ref index field.

RTP Hint Tracks

RTP hint tracks contain information that allows a streaming server to create RTP streams from a QuickTime movie, without requiring the server to know anything about the media type, compression, or payload format.



In RTP, each media stream, such as an audio or video track, is sent as a separate RTP stream. Consequently, each media track in the movie has an associated RTP hint track containing the data necessary to packetize it for RTP transport, and each hint track contains a track reference back to its associated media track.

Media tracks that do not have an associated RTP hint track cannot be streamed over RTP, and should be ignored by RTP streaming servers.

It is possible for a media track to have more than one associated hint track. The hint track contains information such as the packet size and time scale in the hint track's sample description. This minimizes the runtime server load, but in order to support multiple packet sizes, it is necessary to have multiple RTP hint tracks for each media track, each with different a packet size. A similar mechanism could be used to provide hint tracks for multiple protocols in the future.

It is also possible for a single hint track to refer to more than one media stream. For example, audio and video MPEG elementary streams could be multiplexed into a single systems stream RTP payload format, and a single hint track would contain the necessary information to combine both elementary streams into a single series of RTP packets.

This is the exception rather than the rule, however. In general, multiplexing is achieved by using IP's port-level multiplexing, not by interleaving the data from multiple streams into a single RTP session

The hint track is related to each base media track by a track reference declaration. The sample description for RTP declares the maximum packet size which this hint track will generate. Partial session description (SDP) information is stored in the track's user data atom.

Hint Sample Data Format

The sample description atom ('stsd') contains information about the hint track samples. It specifies the data format (note that currently only RTP data format is defined) and which data reference to use (if more than one is defined) to locate the hint track sample data. It also contains some general information about this hint track, such as the hint track version number, the maximum packet size allowed by this hint track, and the RTP time scale. It may contain additional information, such as the random offsets to add to the RTP time stamp and sequence number.

Media Data Atom Types

The sample description atom can contain a table of sample descriptions to accommodate media that are encoded in multiple formats, but a hint track can be expected to have a single sample description at this time.

The sample description for hint tracks is defined in Table 3-9.

Table 3-9 Hint track sample description

Field	Bytes
Size	4
Data format	4
Reserved	6
Data reference index	2
Hint track version	2
Last compatible hint track version	2
Max packet size	4
Additional data table	variable
Field descriptions	
Size	A 32-bit integer specifying the size of this sample description in bytes.
Data format	A four-character code indicating the data format of the hint track samples. Only 'rtp ' is currently defined. Note that the fourth character in 'rtp ' is an ASCII blank space (hex 20). Do not attempt to packetize data whose format you do not recognize.
Reserved	Six bytes that must be set to 0.
Data reference index	This field indirectly specifies where to find the hint track sample data. The data reference is a file or resource specified by the data reference atom ('dref') inside the data information atom ('dinf') of the hint track. The data information atom can contain a table of data references, and the data reference index is a 16-bit integer that tells you which entry in that table should be used. Normally, the hint

Media Data Atom Types

track has a single data reference, and this index entry is set to 0.

Hint track version

A 16-bit unsigned integer indicating the version of the hint track specification. This is currently set to 1.

Last compatible hint track version

A 16-bit unsigned integer indicating the oldest hint track version with which this hint track is backward-compatible. If your application understands the hint track version specified by this field, it can work with this hint track.

Max packet size

A 32-bit integer indicating the packet size limit, in bytes, used when creating this hint track. The largest packet generated by this hint track will be no larger than this limit.

Additional data table

A table of variable length containing additional information. Additional information is formatted as a series of tagged entries.

This field always contains a tagged entry indicating the RTP time scale for RTP data. All other tagged entries are optional.

Three data tags are currently defined for RTP data. One tag is defined for use with any type of data. You can create additional tags. Tags are identified using four-character codes. Tags using all lowercase letters are reserved by Apple. Ignore any tagged data you do not understand.

Table entries are structured like atoms. The structure of table entries is shown in Table 3-10.

Table 3-10 The structure of table entries

Field	Format	Bytes
Entry length	32-bit integer	4
Data tag	4-char code	4
Data	Variable	Entry length - 8

Tagged entries for the 'rtp ' data format are defined as follows:

'tims'	A 32-bit integer specifying the RTP time scale. This entry is required for RTP data.
'tsro'	A 32-bit integer specifying the offset to add to the stored time stamp when sending RTP packets. If this entry is not present, a random offset should be used, as specified by the IETF. If this entry is 0, use an offset of 0 (no offset).
'snro'	A 32-bit integer specifying the offset to add to the sequence number when sending RTP packets. If this entry is not present, a random offset should be used, as specified by the IETF. If this entry is 0, use an offset of 0 (no offset).

Packetization Hint Sample Data for Data Format 'rtp '

This section describes the sample data for the 'rtp ' format. The 'rtp ' format assumes that the server is sending data using Real-Time Transport Protocol (RTP). This format also assumes that the server “knows” about RTP headers but does not require that the server know anything about specific media headers, including media headers defined in various IETF drafts.

Each sample in the hint track will generate one or more RTP packets. Each entry in the sample data table in a hint track sample corresponds to a single RTP packet. Samples in the hint track may or may not correspond exactly to samples in the media track. Data in the hint track sample is byte aligned, but not 32-bit aligned.

The RTP timestamps of all packets in a hint sample are the same as the hint sample time. In other words, packets that do not have the same RTP timestamp cannot be placed in the same hint sample.

The RTP hint track time scale should be reasonably chosen so that there is adequate spacing between samples (as well as adequate spacing between transmission times for packets within a sample).

The packetization hint sample data contains the following data elements.

Packetization hint sample data	Bytes
Entry count	2
Reserved	2
Packet entry table	variable
Additional data	variable

Field descriptions

Entry count	A 16-bit unsigned integer indicating the number of packet entries in the table. Each entry in the table corresponds to a packet. Multiple entries in a single sample indicate that the media sample had to be split into multiple packets. A sample with an entry count of 0 is reserved and, if encountered, must be skipped.
Reserved	Two bytes that must be set to 0.
Packet entry table	A variable length table containing packet entries. Packet entries are defined below.
Additional data	A variable length field containing data pointed to by the entries in the data table.

The packet entry contains the following data elements.

Packet entry	Bytes
Relative packet transmission time	4
RTP header info	2
RTP sequence number	2
Flags	2
Entry count	2
Extra information TLVs	0 or variable
Data table	variable

Field descriptions

Relative packet transmission time

A 32-bit signed integer value, indicating the time, in the hint track's time scale, to send this packet relative to the hint sample's actual time. Negative values mean that the packet will be sent earlier than real time, which is useful for smoothing the data rate. Positive values are useful for repeating packets at later times. Within each hint sample track, each packet time stamp must be non-decreasing.

RTP header info A 16-bit integer specifying various values to be set in the RTP header. The bits of the field are defined as follows.

0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5
reserved		P	X	reserved				M	payload type						

The RTP header information field contains the following elements.

Field	Bit#	Description
P	2	A 1-bit number corresponding to the padding (P) bit in the RTP header. This bit should probably not be set, since a server that needs different packet padding would need to un-pad and re-pad the packet itself.
X	3	A 1-bit number corresponding to the extension (X) bit in the RTP header. This bit should probably not be set, since a server that needs to send its own RTP extension would either not be able to, or would be forced to replace any extensions from the hint track.
M	8	A 1-bit number corresponding to the marker (M) bit in the RTP header.
Payload type	9-15	A 7-bit number corresponding to the payload type (PT) field of the RTP header.

All undefined bits are reserved and must be set to zero. Note that the location of the defined bits are in the same bit location as in the RTP header.

RTP sequence number

A 16-bit integer specifying the RTP sequence number for this packet. The RTP server adds a random offset to this

Media Data Atom Types

sequence number before transmitting the packet. This field allows re-transmission of packets—for example, the same packet can be assembled with the same sequence number and a different (later) packet transmission time. A text sample with a duration of 5 minutes can be retransmitted every 10 seconds, so that clients that miss the original sample transmission (perhaps they started playing the movie in the middle) will be refreshed after a maximum of 10 seconds.

Flags A 16-bit field indicating certain attributes for this packet. Defined bits are:

0	1	...	...	12	13	14	15
reserved				X	B	R	

The RTP header information field contains the following elements.

Field	Bit#	Description
X	13	A 1-bit number indicating that this packet contains an Extra information TLV data table.
B	14	A 1-bit number indicating that this packet contains data that is part of a b-frame. A server that is having difficulty being able to send all the packets in real time may discard packets that have this bit set, until it catches up with the clock.
R	15	A 1-bit number indicating that this is a repeat packet: the data has been defined in a previous packet. A server may choose to skip repeat packets to help it catch up when it is behind in its transmission of packets. All repeated packets for a given packet must live in the same hint sample.

All undefined bits are reserved and must be set to 0.

Entry count A 16-bit unsigned integer specifying the number of entries in the data table.

Extra information TLVs

The extra information TLVs are only present if and only if the X bit is set in the flags field above. This provides a way

of extending the hint track format without changing the version, while allowing backward compatibility.

Extra information TLVs	Bytes
Extra information size	4
TLV size	4
TLV type	4
TLV data	Padded to 4-byte boundary $(\text{int}(\text{TLV Size} - 8 + 3) / 4 * 4)$
TLV size	4
TLV type	4
TLV data	Padded to 4-byte boundary $(\text{int}(\text{TLV Size} - 8 + 3) / 4 * 4)$
TLV size and so forth	...
Extra information size	A 32-bit number that is the total size of all extra information TLVs in this packet, including the 4 bytes used for this field. An empty Extra information TLVs table would just be the Extra information size, having the value 4. (In this case, it would be more efficient simply to not set the X bit and save 4 bytes just to represent the empty table.)
TLV size	A 32-bit number that is the total size of this one TLV entry, including 4 bytes for the size, 4 bytes for the type, and any data bytes, but not including padding required to align to the next 4 byte boundary.
TLV type	A 32-bit tag (a four-character <code>OSType</code>) identifying the TLV. Servers <i>must</i> ignore TLV types that they do not recognize. Note that TLV types containing all lowercase letters are reserved by Apple Computer.
TLV data	The data for the TLV.

In order to support MPEG (and other data types) whose RTP timestamp is not monotonically increasing and directly calculated from the sample timestamp, the following TLV type is defined:

Size	Type	Data Description
12	'rtpp'	A signed 32-bit integer, which is to be added to the RTP timestamp, which is derived from the hint sample timestamp.
Data table		A table that defines the data to be put in the payload portion of the RTP packet. This table defines various places the data can be retrieved.

Data table entry	Bytes
Data source	1
Data	15

The data source field of the entry table indicates how the other 15 bytes of the entry are to be interpreted. Values of 0 through 4 are defined. The various data table formats are defined below.

Although there are various schemes, note that the entries in the various schemes are the same size, 16 bytes long.

No-Op Data Mode

The data table entry has the following format for no-op mode.

0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5
data source = 0								ignored							
ignored															


Field descriptions

Data source = 0	A value of 0 indicates that this data table entry is to be ignored.
-----------------	---

Immediate Data Mode

The data table entry has the following format for immediate mode.

Media Data Atom Types



Field descriptions

- Data source = 1


A value of one indicates that the data is to be immediately taken from the bytes of data that follow.
- Immediate length

An 8-bit integer indicating the number of bytes to take from the data that follows. Legal values range from 0 to 14.
- Immediate data

14 bytes of data to place into the payload portion of the packet. Only the first number of bytes indicated by the immediate length field are used.

Sample Mode

The data table entry has the following format for sample mode.



Field descriptions

Data source = 2	A value of two indicates that the data is to be taken from a track's sample data.
Track ref index	A value that indicates which track the sample data will come from. A value of zero means that there is exactly one media track referenced, so use that. Values from 1 to 127 are indexes into the hint track reference atom entries, indicating which original media track the sample is to be read from. A value of -1 means the hint track itself, that is, get the sample from the same track as the hint sample you are currently parsing.
Length	A 16-bit integer specifying the number of bytes in the sample to copy.
Sample number	A 32-bit integer specifying sample number of the track.
Offset	A 32-bit integer specifying the offset from the start of the sample from which to start copying. If you are referencing samples in the hint track, this will generally points into the Additional Data area.
Bytes per compression block	A 16-bit unsigned integer specifying the number of bytes that results from compressing the number of samples in the Samples per compression block field. A value of 0 is equivalent to a value of 1.
Samples per compression block	A 16-bit unsigned integer specifying the uncompressed samples per compression block. A value of 0 is equivalent to a value of 1.

If the bytes per compression block and /or the samples per compression block is greater than 1, than this ratio is used to translate a sample number into an actual byte offset.

This ratio mode is typically used for compressed audio tracks. Note that for QuickTime sound tracks, the bytes per compression block also factors in the number of sound channels in that stream, so a QuickTime stereo sound stream's BPCB would be twice that of a mono stream of the same sound format.

$$(CB = NS * BPCB / SPCB)$$

Media Data Atom Types

where CB = compressed bytes, NS = number of samples, BPCB = bytes per compression block, and SPCB = samples per compression block.

An example:

A GSM compression block is typically 160 samples packed into 33 bytes.

So, BPCB = 33 and SPCB = 160.

The hint sample requests 33 bytes of data starting at the 161st media sample. Assume that the first QuickTime chunk contains at least 320 samples. So after determining that this data will come from chunk 1, and where chunk 1 starts, you must use this ratio to adjust the offset into the file where the requested samples will be found:

```
chunk_number = 1; /* calculated by walking the sample-to-chunk atom */
first_sample_in_this_chunk = 1; /* also calculated from that atom */
chunk_offset = chunk_offsets[chunk_number]; /* from the stco atom */
data_offset = (sample_number - first_sample_in_this_chunk) * BPCB / SPCB;
read_from_file(chunk_offset + data_offset, length); /* read our data */
```

Sample Description Mode

The data table entry has the following format for sample description mode.

0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5
data source = 3								track ref index							
length															
sample description index															
offset															
reserved															

Field descriptions

Data source = 3 A value of three indicates that the data is to be taken from the media track's sample description table.

Track ref index	A value that indicates which track the sample description will come from. A value of 0 means that there is exactly one hint track reference, so use that. Values from 1 to 127 are indexes into the hint track reference atom entries, indicating which original media track the sample is to be read from. A value of -1 means the hint track itself, that is, get the sample description from the same track as the hint sample you are currently parsing.
Length	A 16-bit integer specifying the number of bytes to copy.
Sample description index	A 32-bit integer specifying the index into the media's sample description table.
Offset	A 32-bit integer specifying the offset from the start of the sample description from which to start copying.
Reserved	Four bytes that must be set to 0.
Field descriptions	
Additional data	A variable length field containing data pointed to by hint track sample mode entries in the data table.

VR Media

This section describes the QuickTime VR world and node information atom containers, which can be obtained by calling the QuickTime VR Manager routines `QTVRGetVRWorld` and `QTVRGetNodeInfo`. Those routines, as well as a complete discussion of QuickTime VR and how your application can create QuickTime VR movies, are described in detail in *Programming With QuickTime VR*, available at <http://developer.apple.com/techpubs/quicktime/quicktime.html>.

Many atom types contained in the VR world and node information atom containers are unique within their container. For example, each has a single header atom. Most parent atoms within an atom container are unique as well, such as the node parent atom in the VR world atom container or the hot spot parent atom in the node information atom container. For these one-time-only atoms, the atom ID is always set to 1. Unless otherwise mentioned in the descriptions of the atoms that follow, assume that the atom ID is 1.

Note that many atom structures contain two version fields, `majorVersion` and `minorVersion`. The values of these fields correspond to the constants `kQTVRMajorVersion` and `kQTVRMinorVersion` found in the header file `QuickTimeVRFormat.h`. For QuickTime 2.0 files, these values are 2 and 0.

QuickTime provides a number of routines for both creating and accessing atom containers.

Some of the leaf atoms within the VR world and node information atom containers contain fields that specify the ID of **string atoms** that are siblings of the leaf atom. For example, the VR world header atom contains a field for the name of the scene. The string atom is a leaf atom whose atom type is `kQTVRStringAtomType` ('vrsg'). Its atom ID is that specified by the referring leaf atom.

A string atom contains a string. The structure of a string atom is defined by the `QTVRStringAtom` data type:

```
typedef struct QTVRStringAtom {
    UInt16                stringUsage;
    UInt16                stringLength;
    unsigned char          theString[4];
} QTVRStringAtom, *QTVRStringAtomPtr;
```

Field descriptions

<code>stringUsage</code>	The string usage. This field is unused.
<code>stringLength</code>	The length, in bytes, of the string.
<code>theString</code>	The string. The string atom structure is extended to hold this string.

Each string atom may also have a sibling leaf atom called the **string encoding atom**. The string encoding atom's atom type is `kQTVRStringEncodingAtomType` ('vrse'). Its atom ID is the same as that of the corresponding string atom. The string encoding atom contains a single variable, `TextEncoding`, a `UInt32`, as defined in the header file `TextCommon.h`. The value of `TextEncoding` is handed, along with the string, to the routine `QTTextToNativeText` for conversion for display on the current machine. The routine `QTTextToNativeText` is found in the header file `Movies.h`.

Note

The header file `TextCommon.h` contains constants and routines for generating and handling text encodings. ♦


VR World Atom Container

The VR world atom container (VR world for short) includes such information as the name for the entire scene, the default node ID, and default imaging properties, as well as a list of the nodes contained in the QTVR track.

A VR world can also contain custom scene information. QuickTime VR ignores any atom types that it doesn't recognize, but you can extract those atoms from the VR world using standard QuickTime atom functions.

The structure of the VR world atom container is shown in Figure 3-6. The component atoms are defined and their structures are shown in the sections that follow.

Figure 3-6 Structure of the VR world atom container



VR World Header Atom Structure

The VR world header atom is a leaf atom. Its atom type is `kQTVRWorldHeaderAtomType ('vrsc')`. It contains the name of the scene and the default node ID to be used when the file is first opened as well as fields reserved for future use.

The structure of a VR world header atom is defined by the `QTVRWorldHeaderAtom` data type:

Media Data Atom Types

```
typedef struct QTVRWorldHeaderAtom {
    UInt16          majorVersion;
    UInt16          minorVersion;
    QTAtomID        nameAtomID;
    UInt32          defaultNodeID;
    UInt32          vrWorldFlags;
    UInt32          reserved1;
    UInt32          reserved2;
} QTVRWorldHeaderAtom, *QTVRWorldHeaderAtomPtr;
```

Field descriptions

majorVersion	The major version number of the file format.
minorVersion	The minor version number of the file format.
nameAtomID	The ID of the string atom that contains the name of the scene. That atom should be a sibling of the VR world header atom. The value of this field is 0 if no name string atom exists.
defaultNodeID	The ID of the default node (that is, the node to be displayed when the file is first opened).
vrWorldFlags	A set of flags for the VR world. This field is unused.
reserved1	Reserved. This field must be 0.
reserved2	Reserved. This field must be 0.

Imaging Parent Atom

The imaging parent atom is the parent atom of one or more node-specific imaging atoms. Its atom type is `kQTVRImagingParentAtomType ('imgp')`. Only panoramas have an imaging atom defined.

Panorama-Imaging Atom

A panorama-imaging atom describes the default imaging characteristics for all the panoramic nodes in a scene. This atom overrides QuickTime VR's own defaults.

The panorama-imaging atom has an atom type of `kQTVRPanoImagingAtomType ('impr')`. Generally, there is one panorama-imaging atom for each imaging mode, so the atom ID, while it must be unique for each atom, is ignored. QuickTime VR iterates through all the panorama-imaging atoms.

Media Data Atom Types

The structure of a panorama-imaging atom is defined by the `QTVRPanoImagingAtom` data type:

```
typedef struct QTVRPanoImagingAtom {
    UInt16          majorVersion;
    UInt16          minorVersion;
    UInt32          imagingMode;
    UInt32          imagingValidFlags;
    UInt32          correction;
    UInt32          quality;
    UInt32          directDraw;
    UInt32          imagingProperties[6];
    UInt32          reserved1;
    UInt32          reserved2;
} QTVRPanoImagingAtom, *VRPanoImagingAtomPtr;
```

Field descriptions

<code>majorVersion</code>	The major version number of the file format.
<code>minorVersion</code>	The minor version number of the file format.
<code>imagingMode</code>	The imaging mode to which the default values apply. Only <code>kQTVRStatic</code> and <code>kQTVRMotion</code> are allowed here.
<code>imagingValidFlags</code>	A set of flags that indicate which imaging property fields in this structure are valid.
<code>correction</code>	The default correction mode for panoramic nodes. This can be either <code>kQTVRNoCorrection</code> , <code>kQTVRPartialCorrection</code> , or <code>kQTVRFullCorrection</code> .
<code>quality</code>	The default imaging quality for panoramic nodes.
<code>directDraw</code>	The default direct-drawing property for panoramic nodes. This can be <code>true</code> or <code>false</code> .
<code>imagingProperties</code>	Reserved for future panorama-imaging properties.
<code>reserved1</code>	Reserved. This field must be 0.
<code>reserved2</code>	Reserved. This field must be 0.

Imaging Property Valid Flags

The `imagingValidFlags` field in the panorama-imaging atom structure specifies which imaging property fields in that structure are valid. You can use these bit flags to specify a value for that field:

Media Data Atom Types

```
enum {
    kQTVRValidCorrection          = 1 << 0,
    kQTVRValidQuality            = 1 << 1,
    kQTVRValidDirectDraw        = 1 << 2,
    kQTVRValidFirstExtraProperty = 1 << 3
};
```

Constant descriptions

`kQTVRValidCorrection`

If this bit is set, the `correction` field holds a default correction mode.

`kQTVRValidQuality` If this bit is set, the `quality` field holds a default imaging quality.

`kQTVRValidDirectDraw`

If this bit is set, the `directDraw` field holds a default direct-drawing property.

`kQTVRValidFirstExtraProperty`

If this bit is set, the first element in the array in the `imagingProperties` field holds a default imaging property. As new imaging properties are added, they will be stored in this array.

Node Parent Atom

The node parent atom is the parent of one or more node ID atoms. The atom type of the node parent atom is `kQTVRNodeParentAtomType ('vrnp')` and the atom type of the each node ID atom is `kQTVRNodeIDAtomType ('vrni')`.

There is one node ID atom for each node in the file. The atom ID of the node ID atom is the node ID of the node. The node ID atom is the parent of the node location atom. The node location atom is the only child atom defined for the node ID atom. Its atom type is `kQTVRNodeLocationAtomType ('nloc')`.

Node Location Atom Structure

The node location atom is the only child atom defined for the node ID atom. Its atom type is `kQTVRNodeLocationAtomType ('nloc')`. A node location atom describes the type of a node and its location.

The structure of a node location atom is defined by the `QTVRNodeLocationAtom` data type:

Media Data Atom Types

```
typedef struct QTVRNodeLocationAtom {
    UInt16          majorVersion;
    UInt16          minorVersion;
    OSType          nodeType;
    UInt32          locationFlags;
    UInt32          locationData;
    UInt32          reserved1;
    UInt32          reserved2;
} QTVRNodeLocationAtom, *QTVRNodeLocationAtomPtr;
```

Field descriptions

<code>majorVersion</code>	The major version number of the file format.
<code>minorVersion</code>	The minor version number of the file format.
<code>nodeType</code>	The node type. This field should contain either <code>kQTVRPanoramaType</code> or <code>kQTVRObjectType</code> .
<code>locationFlags</code>	The location flags. This field must contain the value <code>kQTVRSameFile</code> , indicating that the node is to be found in the current file. In future, these flags may indicate that the node is in a different file or at some URL location.
<code>locationData</code>	The location of the node data. When the <code>locationFlags</code> field is <code>kQTVRSameFile</code> , this field should be 0. The nodes are found in the file in the same order that they are found in the node list.
<code>reserved1</code>	Reserved. This field must be 0.
<code>reserved2</code>	Reserved. This field must be 0.

Custom Cursor Atoms

The hot spot information atom, discussed in “Hot Spot Information Atom” (page 214), allows you to indicate custom cursor IDs for particular hot spots that replace the default cursors used by QuickTime VR. QuickTime VR allows you to store your custom cursors in the VR world of the movie file.

Note

If you’re using the Mac OS, you could store your custom cursors in the resource fork of the movie file. However, this would not work on any other platform (such as Windows), so storing cursors in the resource fork of the movie file is not recommended. ♦

Media Data Atom Types


The cursor parent atom is the parent of all of the custom cursor atoms stored in the VR world. Its atom type is `kQTVRCursorParentAtomType ('vrCP')`. The child atoms of the cursor parent are either cursor atoms or color cursor atoms. Their atom types are `kQTVRCursorAtomType ('CURS')` and `kQTVRColorCursorAtomType ('crsr')`. These atoms are stored exactly as cursors or color cursors would be stored as a resource.

Node Information Atom Container

The node information atom container includes general information about the node such as the node's type, ID, and name. The node information atom container also contains the list of hot spot atoms for the node. A QuickTime VR movie contains one node information atom container for each node in the file. The routine `QTVRGetNodeInfo` allows you to obtain the node information atom container for the current node or for any other node in the movie.

Figure 3-7 shows the structure of the node information atom container.

Figure 3-7 Structure of the node information atom container



Node Header Atom Structure

A node header atom is a leaf atom that describes the type and ID of a node, as well as other information about the node. Its atom type is `kQTVRNodeHeaderAtomType ('ndhd')`.

The structure of a node header atom is defined by the `QTVRNodeHeaderAtom` data type:

Media Data Atom Types

```
typedef struct QTVRNodeHeaderAtom {
    UInt16          majorVersion;
    UInt16          minorVersion;
    OSType          nodeType;
    QTAtomID        nodeID;
    QTAtomID        nameAtomID;
    QTAtomID        commentAtomID;
    UInt32          reserved1;
    UInt32          reserved2;
} QTVRNodeHeaderAtom, *VRNodeHeaderAtomPtr;
```

Field descriptions

majorVersion	The major version number of the file format.
minorVersion	The minor version number of the file format.
nodeType	The node type. This field should contain either <code>kQTVRPanoramaType</code> or <code>kQTVRObjectType</code> .
nodeID	The node ID.
nameAtomID	The ID of the string atom that contains the name of the node. This atom should be a sibling of the node header atom. The value of this field is 0 if no name string atom exists.
commentAtomID	The ID of the string atom that contains a comment for the node. This atom should be a sibling of the node header atom. The value of this field is 0 if no comment string atom exists.
reserved1	Reserved. This field must be 0.
reserved2	Reserved. This field must be 0.

Hot Spot Parent Atom

The hot spot parent atom is the parent for all hot spot atoms for the node. The atom type of the hot spot parent atom is `kQTVRHotSpotParentAtomType ('hspa')` and the atom type of the each hot spot atom is `kQTVRHotSpotAtomType ('hots')`. The atom ID of each hot spot atom is the hot spot ID for the corresponding hot spot. The hot spot ID is determined by its color index value as it is stored in the hot spot image track.

The hot spot track is an 8-bit video track that contains color information that indicates hot spots. For more information, refer to *Programming With QuickTime VR*.

Each hot spot atom is the parent of a number of atoms that contain information about each hot spot.

Hot Spot Information Atom

The hot spot information atom contains general information about a hot spot. Its atom type is `kQTVRHotSpotInfoAtomType ('hsin')`. Every hot spot atom should have a hot spot information atom as a child.

The structure of a hot spot information atom is defined by the `QTVRHotSpotInfoAtom` data type:

```
typedef struct QTVRHotSpotInfoAtom {
    UInt16          majorVersion;
    UInt16          minorVersion;
    OSType          hotSpotType;
    QTAtomID        nameAtomID;
    QTAtomID        commentAtomID;
    SInt32          cursorID[3];
    Float32         bestPan;
    Float32         bestTilt;
    Float32         bestFOV;
    FloatPoint      bestViewCenter;
    Rect            hotSpotRect;
    UInt32          flags;
    UInt32          reserved1;
    UInt32          reserved2;
} QTVRHotSpotInfoAtom, *QTVRHotSpotInfoAtomPtr;
```

Field descriptions

<code>majorVersion</code>	The major version number of the file format.
<code>minorVersion</code>	The minor version number of the file format.
<code>hotSpotType</code>	The hot spot type. This type specifies which other information atoms—if any—are siblings to this one. QuickTime VR recognizes three types: <code>kQTVRHotSpotLinkType</code> , <code>kQTVRHotSpotURLType</code> , and <code>kQTVRHotSpotUndefinedType</code> .
<code>nameAtomID</code>	The ID of the string atom that contains the name of the hot spot. This atom should be a sibling of the hot spot information atom. This string is displayed in the

Media Data Atom Types

	QuickTime VR controller bar when the mouse is moved over the hot spot.
<code>commentAtomID</code>	The ID of the string atom that contains a comment for the hot spot. This atom should be a sibling of the hot spot information atom. The value of this field is 0 if no comment string atom exists.
<code>cursorID</code>	An array of three IDs for custom hot spot cursors (that is, cursors that override the default hot spot cursors provided by QuickTime VR). The first ID (<code>cursorID[0]</code>) specifies the cursor that is displayed when it is in the hot spot. The second ID (<code>cursorID[1]</code>) specifies the cursor that is displayed when it is in the hot spot and the mouse button is down. The third ID (<code>cursorID[2]</code>) specifies the cursor that is displayed when it is in the hot spot and the mouse button is released. To retain the default cursor for any of these operations, set the corresponding cursor ID to 0. Custom cursors should be stored in the VR world atom container, as described in “VR World Atom Container” (page 205).
<code>bestPan</code>	The best pan angle for viewing this hot spot.
<code>bestTilt</code>	The best tilt angle for viewing this hot spot.
<code>bestFOV</code>	The best field of view for viewing this hot spot.
<code>bestViewCenter</code>	The best view center for viewing this hot spot; applies only to object nodes.
<code>hotSpotRect</code>	The boundary box for this hot spot, specified as the number of pixels in full panoramic space. This field is valid only for panoramic nodes.
<code>flags</code>	A set of hot spot flags. This field is unused.
<code>reserved1</code>	Reserved. This field must be 0.
<code>reserved2</code>	Reserved. This field must be 0.

Note

In QuickTime VR movie files, all angular values are stored as 32-bit floating-point values that specify degrees. In addition, all floating-point values conform to the IEEE Standard 754 for binary floating-point arithmetic, in big-endian format. ♦

Specific Information Atoms

Depending on the value of the `hotSpotType` field in the hot spot info atom there may also be a type specific information atom. The atom type of the type-specific atom is the hot spot type.

Link Hot Spot Atom

The link hot spot atom specifies information for hot spots of type `kQTVRHotSpotLinkType ('link')`. Its atom type is thus `'link'`. The link hot spot atom contains specific information about a link hot spot.

The structure of a link hot spot atom is defined by the `QTVRLinkHotSpotAtom` data type:

```
typedef struct VRLinkHotSpotAtom {
    UInt16          majorVersion;
    UInt16          minorVersion;
    UInt32          toNodeID;
    UInt32          fromValidFlags;
    Float32         fromPan;
    Float32         fromTilt;
    Float32         fromFOV;
    FloatPoint      fromViewCenter;
    UInt32          toValidFlags;
    Float32         toPan;
    Float32         toTilt;
    Float32         toFOV;
    FloatPoint      toViewCenter;
    Float32         distance;
    UInt32          flags;
    UInt32          reserved1;
    UInt32          reserved2;
} QTVRLinkHotSpotAtom, *VRLinkHotSpotAtomPtr;
```

Field descriptions

<code>majorVersion</code>	The major version number of the file format.
<code>minorVersion</code>	The minor version number of the file format.
<code>toNodeID</code>	The ID of the destination node (that is, the node to which this hot spot is linked).

Media Data Atom Types

<code>fromValidFlags</code>	A set of flags that indicate which source node view settings are valid.
<code>fromPan</code>	The preferred from-pan angle at the source node (that is, the node containing the hot spot).
<code>fromTilt</code>	The preferred from-tilt angle at the source node.
<code>fromFOV</code>	The preferred from-field of view at the source node.
<code>fromViewCenter</code>	The preferred from-view center at the source node.
<code>toValidFlags</code>	A set of flags that indicate which destination node view settings are valid.
<code>toPan</code>	The pan angle to use when displaying the destination node.
<code>toTilt</code>	The tilt angle to use when displaying the destination node.
<code>toFOV</code>	The field of view to use when displaying the destination node.
<code>toViewCenter</code>	The view center to use when displaying the destination node.
<code>distance</code>	The distance between the source node and the destination node.
<code>flags</code>	A set of link hot spot flags. This field is unused and should be set to 0.
<code>reserved1</code>	Reserved. This field must be 0.
<code>reserved2</code>	Reserved. This field must be 0.

Certain fields in the link hot spot atom are not used by QuickTime VR. The `fromValidFlags` field is generally set to 0 and the other `from` fields are not used. However, these fields could be quite useful if you have created a transition movie from one node to another. The `from` angles can be used to swing the current view of the source node to align with the first frame of the transition movie. The `distance` field is intended for use with 3D applications, but is also not used by QuickTime VR.

Link Hot Spot Valid Flags

The `toValidFlags` field in the link hot spot atom structure specifies which view settings are to be used when moving to a destination node from a hot spot. You can use these bit flags to specify a value for that field:

```
enum {
    kQTVRValidPan           = 1 << 0,
    kQTVRValidTilt         = 1 << 1,
```

Media Data Atom Types

```

    kQTVRValidFOV                = 1 << 2,
    kQTVRValidViewCenter         = 1 << 3
};

```

Constant descriptions

kQTVRValidPan	If this bit is set, the destination pan angle is used.
kQTVRValidTilt	If this bit is set, the destination tilt angle is used.
kQTVRValidFOV	If this bit is set, the destination field of view is used.
kQTVRValidViewCenter	If this bit is set, the destination view center is used.

URL Hot Spot Atom

The URL hot spot atom has an atom type of `kQTVRHotSpotURLType ('url ')`. The URL hot spot atom contains a URL string for a particular Web location (for example, `http://quicktimevr.apple.com`). QuickTime VR automatically links to this URL when the hot spot is clicked.

Support for Wired Actions

Certain actions on a QuickTime VR movie can trigger wired actions if the appropriate event handler atoms have been added to the file. This section discusses what atoms must be included in the QuickTime VR file to support wired actions.

As with sprite tracks, the presence of a certain atom in the media property atom container of the QTVR track enables the handling of wired actions. This atom is of type `kSpriteTrackPropertyHasActions`, which has a single Boolean value that must be set to `true`.

When certain events occur and the appropriate event handler atom is found in the QTVR file, then that atom is passed to QuickTime to perform any actions specified in the atom. The event handler atoms themselves must be added to the node information atom container in the QTVR track. There are two types of event handlers for QTVR nodes: global and hot spot specific. The currently supported global event handlers are `kQTEventFrameLoaded` and `kQTEventIdle`. The event handler atoms for these are located at the root level of the node information atom container. A global event handler atom's type is set to the event type and its ID is set to 1.

Media Data Atom Types

Hot spot–specific event handler atoms are located in the specific hot spot atom as a sibling to the hot spot info atom. For these atoms, the atom type is always `kQTEventType` and the ID is the event type. Supported hot spot–specific event types are `kQTEventMouseClicked`, `kQTEventMouseClickedEnd`, `kQTEventMouseClickedEndTriggerButton`, and `kQTEventMouseEnter`, `kQTEventMouseExit`.

The specific actions that cause these events to be generated are described as follows:

`kQTEventFrameLoaded ('fram')`

Generated when a node is entered, before any application-installed entering-node procedure is called (this event processing is considered part of the node setup that occurs before the application’s routine is called).

`kQTEventIdle ('idle')`

Generated every *n* ticks, where *n* is defined by the contents of the `kSpriteTrackPropertyQTIdleEventsFrequency` atom (`SInt32`) in the media property atom container. When appropriate, this event is triggered before any normal idle processing occurs for the QuickTime VR movie.

`kQTEventMouseClicked ('clik')`

Generated when the mouse goes down over a hot spot.

`kQTEventMouseClickedEnd ('cend')`

Generated when the mouse goes up after a `kQTEventMouseClicked` is generated, regardless of whether the mouse is still over the hot spot originally clicked. This event occurs prior to QuickTime VR’s normal mouse-up processing.

`kQTEventMouseClickedEndTriggerButton ('trig')`

Generated when a click end triggers a hot spot (using the same criterion as used by QuickTime VR in 2.1 for link/ url hot spot execution). This event occurs prior to QuickTime VR’s normal hot spot–trigger processing.

`kQTEventMouseEnter ('entr')`, `kQTEventMouseExit ('exit')`

These two events are generated when the mouse rolls into or out of a hot spot, respectively. These events occur whether or not the mouse is down and whether or not the movie is being panned. These events occur after any

application-installed `MouseOverHotSpotProc` is called, and will be cancelled if the return value from the application's routine indicates that QuickTimeVR's normal over-hot spot processing should not take place.

QuickTime VR File Format

A QuickTime VR movie is stored on disk in a format known as the **QuickTime VR file format**. Beginning in QuickTime VR 2.0, a QuickTime VR movie could contain one or more nodes. Each node is either a panorama or an object. In addition, a QuickTime VR movie could contain various types of hot spots, including links between any two types of nodes.

IMPORTANT

This section describes the file format supported by version 2.1 of the QuickTime VR Manager. For information on the file format supported by earlier versions of QuickTime VR, see Macintosh Technotes numbers 1035 and 1036. The Macintosh Technotes are available electronically on the *Developer CD Series* and on the Technote website at <http://devworld.apple.com/dev/technotes.shtml> ▲


All QuickTime VR movies contain a single **QTVR track**, a special type of QuickTime track that maintains a list of the nodes in the movie. Each individual sample in a QTVR track contains general information and hot spot information for a particular node.

If a QuickTime VR movie contains any panoramic nodes, that movie also contains a single **panorama track**, and if it contains any object nodes, it also contains a single **object track**. The panorama and object tracks contain information specific to the panoramas or objects in the movie. The actual image data for both panoramas and objects is usually stored in standard QuickTime video tracks, hereafter referred to as **image tracks**. (An image track can also be any type of track that is capable of displaying an image, such as a QuickTime 3D track.) The individual frames in the image track for a panorama make up the diced frames of the original single panoramic image. The frames for the image track of an object represent the many different views of the object. Hot spot image data is stored in parallel video tracks for both panoramas and objects.

Single-Node Panoramic Movies

Figure 3-8 illustrates the basic structure of a single-node panoramic movie. As you can see, every panoramic movie contains at least three tracks: a QTVR track, a panorama track, and a panorama image track.

Figure 3-8 The structure of a single-node panoramic movie file




For a single-node panoramic movie, the QTVR track contains just one sample. There is a corresponding sample in the panorama track, whose time and duration are the same as the time and duration of the sample in the QTVR track. The time base of the movie is used to locate the proper video samples in the panorama image track. For a panoramic movie, the video sample for the first diced frame of a node's panoramic image is located at the same time as the corresponding QTVR and panorama track samples. The total duration of all the video samples is the same as the duration of the corresponding QTVR sample and the panorama sample.

A panoramic movie can contain an optional hot spot image track and any number of standard QuickTime tracks. A panoramic movie can also contain panoramic image tracks with a lower resolution. The video samples in these low-resolution image tracks must be located at the same time and must have the same total duration as the QTVR track. Likewise, the video samples for a hot spot image track, if one exists, must be located at the same time and must have the same total duration as the QTVR track.

Single-Node Object Movies

Figure 3-9 illustrates the basic structure of a single-node object movie. As you can see, every object movie contains at least three tracks: a QTVR track, an object track, and an object image track.

Figure 3-9 The structure of a single-node object movie file



For a single-node object movie, the QTVR track contains just one sample. There is a corresponding sample in the object track, whose time and duration are the same as the time and duration of the sample in the QTVR track. The time base of the movie is used to locate the proper video samples in the object image track.

For an object movie, the frame corresponding to the first row and column in the object image array is located at the same time as the corresponding QTVR and object track samples. The total duration of all the video samples is the same as the duration of the corresponding QTVR sample and the object sample.

In addition to these three required tracks, an object movie can also contain a hot spot image track and any number of standard QuickTime tracks (such as video, sound, and text tracks). A hot spot image track for an object is a QuickTime video track that contains images of colored regions delineating the hot spots; an image in the hot spot image track must be synchronized to match the appropriate image in the object image track. A hot spot image track should be 8 bits deep and can be compressed with any lossless compressor (including temporal compressors). This is also true of panoramas.

Note

To assign a single fixed-position hot spot to all views of an object, you should create a hot spot image track that consists of a single video frame whose duration is the entire node time. ♦

To play a time-based track with the object movie, you must synchronize the sample data of that track to the start and stop times of a view in the object image track. For example, to play a different sound with each view of an object, you might store a sound track in the movie file with each set of sound samples synchronized to play at the same time as the corresponding object's view image. (This technique also works for video samples.) Another way to add sound or video is simply to play a sound or video track during the object's view animation; to do this, you need to add an active track to the object that is equal in duration to the object's row duration.


IMPORTANT

In a QuickTime VR movie file, the panorama image tracks and panorama hot spot tracks must be disabled. For an object, the object image tracks must be enabled and the object hot spot tracks must be disabled. ▲

Multinode Movies

A multinode QuickTime VR movie can contain any number of object and panoramic nodes. Figure 3-10 illustrates the structure of a QuickTime VR movie that contains five nodes (in this case, three panoramic nodes and two object nodes).

Figure 3-10 The structure of a multinode movie file



IMPORTANT

Panoramic tracks and object tracks must never be located at the same time. ▲

QTVR Track

A QTVR track is a special type of QuickTime track that maintains a list of all the nodes in a movie. The media type for a QTVR track is 'qtvvr'. All the media samples in a QTVR track share a common sample description. This sample description contains the VR world atom container. The track contains one media sample for each node in the movie. Each QuickTime VR media sample contains a node information atom container.

QuickTime VR Sample Description Structure

Whereas the QuickTime VR media sample is simply the node information itself, all sample descriptions are required by QuickTime to have a certain structure for the first several bytes. The structure for the QuickTime VR sample description is as follows:

Media Data Atom Types

```
typedef struct QTVRSampleDescription {
    UInt32                size;
    UInt32                type;
    UInt32                reserved1;
    UInt16                reserved2;
    UInt16                dataRefIndex;
    UInt32                data;
} QTVRSampleDescription, *QTVRSampleDescriptionPtr,
**QTVRSampleDescriptionHandle;
```

Field descriptions

size	The size, in bytes, of the sample description header structure, including the VR world atom container contained in the <code>data</code> field.
type	The sample description type. For QuickTime VR movies, this type should be 'qtvvr'.
reserved1	Reserved. This field must be 0.
reserved2	Reserved. This field must be 0.
dataRefIndex	Reserved. This field must be 0.
data	The VR world atom container. The sample description structure is extended to hold this atom container.

Panorama Tracks

A movie's **panorama track** is a track that contains information about the panoramic nodes in a scene. The media type of the panorama track is 'pano'. Each sample in a panorama track corresponds to a single panoramic node. This sample parallels the corresponding sample in the QTVR track. Panorama tracks do not have a sample description (although QuickTime requires that you specify a dummy sample description when you call `AddMediaSample` to add a sample to a panorama track). The sample itself contains an atom container that includes a panorama sample atom and other optional atoms.

Panorama Sample Atom Structure

A panorama sample atom has an atom type of `kQTVRPanoSampleDataAtomType` ('pdatt'). It describes a single panorama, including track reference indexes of the scene and hot spot tracks and information about the default viewing angles and the source panoramic image.

Media Data Atom Types

The structure of a panorama sample atom is defined by the `QTVRPanoSampleAtom` data type:

```
typedef struct QTVRPanoSampleAtom {
    UInt16      majorVersion;
    UInt16      minorVersion;
    UInt32      imageRefTrackIndex;
    UInt32      hotSpotRefTrackIndex;
    Float32     minPan;
    Float32     maxPan;
    Float32     minTilt;
    Float32     maxTilt;
    Float32     minFieldOfView;
    Float32     maxFieldOfView;
    Float32     defaultPan;
    Float32     defaultTilt;
    Float32     defaultFieldOfView;
    UInt32      imageSizeX;
    UInt32      imageSizeY;
    UInt16      imageNumFramesX;
    UInt16      imageNumFramesY;
    UInt32      hotSpotSizeX;
    UInt32      hotSpotSizeY;
    UInt16      hotSpotNumFramesX;
    UInt16      hotSpotNumFramesY;
    UInt32      flags;
    OSType      panoType;
    UInt32      reserved2;
} QTVRPanoSampleAtom, *QTVRPanoSampleAtomPtr;
```

Field descriptions

<code>majorVersion</code>	The major version number of the file format.
<code>minorVersion</code>	The minor version number of the file format.
<code>imageRefTrackIndex</code>	The index of the image track reference. This is the index returned by the <code>AddTrackReference</code> function when the image track is added as a reference to the panorama track. There can be more than one image track for a given panorama track and hence multiple references. (A panorama track might have multiple image tracks if the panoramas have different characteristics, which could occur if the

Media Data Atom Types

panoramas were shot with different size camera lenses.) The value in this field is 0 if there is no corresponding image track.

hotSpotRefTrackIndex

The index of the hot spot track reference.

minPan

The minimum pan angle, in degrees. For a full panorama, the value of this field is usually 0.0.

maxPan

The maximum pan angle, in degrees. For a full panorama, the value of this field is usually 360.0.

minTilt

The minimum tilt angle, in degrees. For a high-resolution panorama, a typical value for this field is -42.5 .

maxTilt

The maximum tilt angle, in degrees. For a high-resolution panorama, a typical value for this field is $+42.5$.

minFieldOfView

The minimum vertical field of view, in degrees. For a high-resolution panorama, a typical value for this field is 5.0. The value in this field is 0 for the default minimum field of view, which is 5 percent of the maximum field of view.

maxFieldOfView

The maximum vertical field of view, in degrees. For a high-resolution panorama, a typical value for this field is 85.0. The value in this field is 0 for the default maximum field of view, which is $\text{maxTilt} - \text{minTilt}$.

defaultPan

The default pan angle, in degrees.

defaultTilt

The default tilt angle, in degrees.

defaultFieldOfView

The default vertical field of view, in degrees.

imageSizeX

The width, in pixels, of the panorama stored in the highest resolution image track.

imageSizeY

The height, in pixels, of the panorama stored in the highest resolution image track.

imageNumFramesX

The number of frames into which the panoramic image is diced horizontally. The width of each frame (which is $\text{imageSizeX} / \text{imageNumFramesX}$) should be divisible by 4.

imageNumFramesY

The number of frames into which the panoramic image is diced vertically. The height of each frame (which is $\text{imageSizeY} / \text{imageNumFramesY}$) should be divisible by 4.

Media Data Atom Types

hotSpotSizeX	The width, in pixels, of the panorama stored in the highest resolution hot spot image track.
hotSpotSizeY	The height, in pixels, of the panorama stored in the highest resolution hot spot image track.
hotSpotNumFramesX	The number of frames into which the panoramic image is diced horizontally for the hot spot image track.
hotSpotNumFramesY	The number of frames into which the panoramic image is diced vertically for the hot spot image track.
flags	A set of panorama flags. <code>kQTVRPanoFlagHorizontal</code> has been superseded by the <code>panoType</code> field. It is only used when the <code>panoType</code> field is <code>nil</code> to indicate a horizontally-oriented cylindrical panorama. <code>kQTVRPanoFlagAlwaysWrap</code> is set if the panorama should wrap horizontally, regardless of whether or not the pan range is 360 degrees. Note that these flags are currently supported only under Mac OS X.
panoType	An <code>OSType</code> describing the type of panorama. Types supported are <code>kQTVRHorizontalCylinder</code> <code>kQTVRVerticalCylinder</code> <code>kQTVRCube</code>
reserved2	Reserved. This field must be 0.

IMPORTANT

A new flag has been added to the `flags` field of the `QTVRPanoSampleAtom` data structure. This flag controls how panoramas wrap horizontally. If `kQTVRPanoFlagAlwaysWrap` is set, then the panorama wraps horizontally, regardless of the number of degrees in the panorama. If the flag is not set, then the panorama wraps only when the panorama range is 360 degrees. This is the default behavior. ▲

The minimum and maximum values in the panorama sample atom describe the physical limits of the panoramic image. QuickTime VR allows you to set further constraints on what portion of the image a user can see by calling the `QTVRSetConstraints` routine. You can also preset image constraints by adding constraint atoms to the panorama sample atom container. The three constraint atom types are `kQTVRPanConstraintAtomType`, `kQTVRTiltConstraintAtomType`, and `kQTVRFOVConstraintAtomType`. Each of these atom types share a common structure defined by the `QTVRAngleRangeAtom` data type:

Media Data Atom Types

```
typedef struct QTVRAngleRangeAtom {
    Float32          minimumAngle;
    Float32          maximumAngle;
} QTVRAngleRangeAtom, *QTVRAngleRangeAtomPtr;
```

Field descriptions

<code>minimumAngle</code>	The minimum angle in the range, in degrees.
<code>maximumAngle</code>	The maximum angle in the range, in degrees.


Panorama Image Track

The actual panoramic image for a panoramic node is contained in a **panorama image track**, which is a standard QuickTime video track. The track reference to this track is stored in the `imageRefTrackIndex` field of the panorama sample atom.

QuickTime VR 2.1 required the original panoramic image to be rotated 90 degrees counterclockwise. This orientation has changed in QuickTime VR 2.2, however, as discussed later in this section.

The rotated image is diced into smaller frames, and each diced frame is then compressed and added to the video track as a video sample, as shown in Figure 3-11. Frames can be compressed using any spatial compressor; however, temporal compression is not allowed for panoramic image tracks.

Original panorama




QuickTime VR 2.2 does *not* require the original panoramic image to be rotated 90 degrees counterclockwise, as was the case in QuickTime VR 2.1. The rotated image is still diced into smaller frames, and each diced frame is then compressed and added to the video track as a video sample, as shown in Figure 3-12.

Figure 3-12 Creating an image track for a panorama, with the image track oriented horizontally

Original panorama




↓ Dice the Image



↓ Add samples to image track

Image track



In QuickTime 3.0, a panorama sample atom (which contains information about a single panorama) contains the `panoType` field, which indicates whether the diced panoramic image is oriented horizontally or vertically.

Cylindrical Panoramas

The primary change to cylindrical panoramas in QuickTime VR 2.2 is that the panorama, as stored in the image track of the movie, can be oriented horizontally. This means that the panorama does *not* need to be rotated 90 degrees counterclockwise, as required previously.

To indicate a horizontal orientation, the field in the `VRPanoSampleAtom` data structure formerly called `reserved1` has been renamed `panoType`. Its type is `OSType`. The `panoType` for a horizontally oriented cylinder is `kQTVRHorizontalCylinder ('hcyl')`, while a vertical cylinder is `kQTVRVerticalCylinder ('vcyl')`. For compatibility with older QuickTime VR files, when the `panoType` field is `nil`, then a cylinder is assumed, with the low order bit of the flags field set to 1 to indicate if the cylinder is horizontal and 0 if the cylinder is vertical.

One consequence of reorienting the panorama horizontally is that, when the panorama is divided into separate tiles, the order of the samples in the file is now the reverse of what it was for vertical cylinders. Since vertical cylinders were rotated 90 degrees counterclockwise, the first tile added to the image track was the right-most tile in the panorama. For unrotated horizontal cylinders, the first tile added to the image track is the left-most tile in the panorama.

Cubic Panoramas



A new type of panorama was introduced in the current version of QuickTime: the cubic panorama. This panorama in its simplest form is represented by six faces of a cube, thus enabling the viewer to see all the way up and all the way down. The file format and the cubic rendering engine actually allow for more complicated representations, such as special types of cubes with elongated sides or cube faces made up of separate tiles. Atoms that describe the orientation of each face allow for these nonstandard representations. If these atoms are not present, then the simplest representation is assumed. The following describes this simplest representation: a cube with six square sides.

Tracks in a cubic movie are laid out as they are for cylindrical panoramas. This includes a QTVR track, a panorama track, and an image track. Optionally, there

may also be a hot spot track and a fast-start preview track. The image, hot spot, and preview tracks are all standard QuickTime video tracks.

Image Tracks in Cubic Nodes

For a cubic node the image track contains six samples that correspond to the six square faces of the cube. The same applies to hot spot and preview tracks. The following diagram shows how the order of samples in the track corresponds to the orientation of the cube faces.



Note that the frames are oriented horizontally. There is no provision for frames that are rotated 90° counterclockwise as there are for cylindrical panoramas.

Panorama Tracks in Cubic Nodes

The media sample for a panorama track contains the pano sample atom container. For cubes, some of the fields in the pano sample data atom have special values, which provide compatibility back to QuickTime VR 2.2. The cubic projection engine ignores these fields. They allow one to view cubic movies in older versions of QuickTime VR using the cylindrical engine, although the view will be somewhat incorrect, and the top and bottom faces will not be visible. The special values are shown in Table 3-11.

Table 3-11 Fields and their special values as represented in the pano sample data atom, providing backward compatibility to QuickTime VR 2.2

Field	Value
imageNumFramesX	4
imageNumFramesY	1
imageSizeX	frame width * 4
imageSizeY	frame height
minPan	0.0
maxPan	360.0
minTilt	-45.0
maxTilt	45.0
minFieldOfView	5.0
maxFieldOfView	90.0
flags	1

A 1 value in the flags field tells QuickTime VR 2.2 that the frames are not rotated. QuickTime VR 2.2 treats this as a four-frame horizontal cylinder. The `panoType` field (formerly `reserved1`) must be set to `kQTVRCube` ('cube') so that QuickTime VR 3.0 can recognize this panorama as a cube.

Since certain viewing fields in the pano sample data atom are being used for backward compatibility, a new atom must be added to indicate the proper viewing parameters for the cubic image. This atom is the cubic view atom (atom type 'cuvw'). The data structure of the cubic view atom is as follows:

```
struct QTVRCubicViewAtom {
    Float32    minPan;
    Float32    maxPan;
    Float32    minTilt;
    Float32    maxTilt;
    Float32    minFieldOfView;
    Float32    maxFieldOfView;

    Float32    defaultPan;
    Float32    defaultTilt;
```

Media Data Atom Types

```

        Float32          defaultFieldOfView;
    };
    typedef struct QTVRCubicViewAtom    QTVRCubicViewAtom;

```

The fields are filled in as desired for the cubic image. This atom is ignored by older versions of QuickTime VR. Typical values for the `min` and `max` fields are shown in Table 3-12.

Table 3-12 Values for min and max fields

Field	Value
<code>minPan</code>	0.0
<code>maxPan</code>	360.0
<code>minTilt</code>	-90.0
<code>maxTilt</code>	90.0
<code>minFieldOfView</code>	5.0
<code>maxFieldOfView</code>	120.0

You add the cubic view atom to the pano sample atom container (after adding the pano sample data atom). Then use `AddMediaSample` to add the atom container to the panorama track.

Nonstandard Cubes

Although the default representation for a cubic panorama is that of six square faces of a cube, it is possible to depart from this standard representation. When doing so, a new atom must be added to the pano sample atom container. The atom type is `'cufa'`. The atom is an array of data structures of type `QTVRCubicFaceData`. Each entry in the array describes one face of whatever polyhedron is being defined. `QTVRCubicFaceData` is defined as follows:

```

struct QTVRCubicFaceData {
    float    orientation[4];
    float    center[2];
    float    aspect;
    float    skew;
};
typedef struct QTVRCubicFaceData    QTVRCubicFaceData;

```

The mathematical explanation of these data structures is beyond the scope of this document but will be described in a separate Apple Technote. Table 3-13 shows what values QuickTime VR uses for the default representation of six square sides.

Table 3-13 Values used for representing six square sides

Orientation				Center		Aspect	Skew	
1	0	0	0	0	0	1	0	# front
$-\sqrt{.5}$	0	$\sqrt{.5}$	0	0	0	1	0	# right
0	0	1	0	0	0	1	0	# back
$\sqrt{.5}$	0	$\sqrt{.5}$	0	0	0	1	0	# left
$\sqrt{.5}$	$\sqrt{.5}$	0	0	0	0	1	0	# top
$-\sqrt{.5}$	$\sqrt{.5}$	0	0	0	0	1	0	# bottom

Hot Spot Image Tracks

When a panorama contains hot spots, the movie file contains a **hot spot image track**, a video track that contains a parallel panorama, with the hot spots designated by colored regions. Each diced frame of the hot spot panoramic image must be compressed with a lossless compressor (such as QuickTime’s graphics compressor). The dimensions of the hot spot panoramic image are usually the same as those of the image track’s panoramic image, but this is not required. The dimensions must, however, have the same aspect ratio as the image track’s panoramic image. A hot spot image track should be 8 bits deep.

Low-Resolution Image Tracks

It’s possible to store one or more low-resolution versions of a panoramic image in a movie file; those versions are called **low-resolution image tracks**. If there is not enough memory at runtime to use the normal image track, QuickTime VR uses a lower resolution image track if one is available. A low-resolution image track contains diced frames just like the higher resolution track, but the reconstructed panoramic image is half the height and half the width of the higher resolution image.

IMPORTANT

The panoramic images in the lower resolution image tracks and the hot spot image tracks, if present, must have the same orientation (horizontal or vertical) as the panorama image track. ▲

Track Reference Entry Structure

Since there are no fields in the pano sample data atom to indicate the presence of low-resolution image tracks, a separate sibling atom must be added to the panorama sample atom container. The track reference array atom contains an array of track reference entry structures that specify information about any low-resolution image tracks contained in a movie. Its atom type is `kQTVRTrackRefArrayAtomType ('tref')`.

A track reference entry structure is defined by the `QTVRTrackRefEntry` data type:

```
typedef struct QTVRTrackRefEntry {
    UInt32          trackRefType;
    UInt16          trackResolution;
    UInt32          trackRefIndex;
} QTVRTrackRefEntry;
```

Field descriptions

<code>trackRefType</code>	The track reference type.
<code>trackResolution</code>	The track resolution.
<code>trackRefIndex</code>	The index of the track reference.

The number of entries in the track reference array atom is determined by dividing the size of the atom by `sizeof (QTVRTrackRefEntry)`.

`kQTVRPreviewTrackRes` is a special value for the `trackResolution` field in the `QTVRTrackRefEntry` structure. This is used to indicate the presence of a special preview image track.

Object Tracks

A movie's **object track** is a track that contains information about the object nodes in a scene. The media type of the object track is `'obje'`. Each sample in an object track corresponds to a single object node in the scene. The samples of the

object track contain information describing the object images stored in the object image track.

These object information samples parallel the corresponding node samples in the QTVR track and are equal in time and duration to a particular object node's image samples in the object's image track as well as the object node's hot spot samples in the object's hot spot track.

Object tracks do not have a sample description (although QuickTime requires that you specify a dummy sample description when you call `AddMediaSample` to add a sample to an object track). The sample itself is an atom container that contains a single object sample atom and other optional atoms.

Object Sample Atom Structure

object sample atom describes a single object, including information about the default viewing angles and the view settings. The structure of an object sample atom is defined by the `QTVRObjectSampleAtom` data type:

```
typedef struct QTVRObjectSampleAtom {
    UInt16          majorVersion;
    UInt16          minorVersion;
    UInt16          movieType;
    UInt16          viewStateCount;
    UInt16          defaultViewState;
    UInt16          mouseDownViewState;
    UInt32          viewDuration;
    UInt32          columns;
    UInt32          rows;
    Float32         mouseMotionScale;
    Float32         minPan;
    Float32         maxPan;
    Float32         defaultPan;
    Float32         minTilt;
    Float32         maxTilt;
    Float32         defaultTilt;
    Float32         minFieldOfView;
    Float32         fieldOfView;
    Float32         defaultFieldOfView;
    Float32         defaultViewCenterH;
    Float32         defaultViewCenterV;
    Float32         viewRate;
```

Media Data Atom Types

```

Float32      frameRate;
UInt32      animationSettings;
UInt32      controlSettings;
} QTVRObjectSampleAtom, *QTVRObjectSampleAtomPtr;

```

Field descriptions

majorVersion	The major version number of the file format.
minorVersion	The minor version number of the file format.
movieType	The movie controller type.
viewStateCount	The number of view states of the object. A view state selects an alternate set of images for an object's views. The value of this field must be positive.
defaultViewState	The 1-based index of the default view state. The default view state image for a given view is displayed when the mouse button is not down.
mouseDownViewState	The 1-based index of the mouse-down view state. The mouse-down view state image for a given view is displayed while the user holds the mouse button down and the cursor is over an object movie.
viewDuration	The total movie duration of all image frames contained in an object's view. In an object that uses a single frame to represent a view, the duration is the image track's sample duration time.
columns	The number of columns in the object image array (that is, the number of horizontal positions or increments in the range defined by the minimum and maximum pan values). The value of this field must be positive.
rows	The number of rows in the object image array (that is, the number of vertical positions or increments in the range defined by the minimum and maximum tilt values). The value of this field must be positive.
mouseMotionScale	The mouse motion scale factor (that is, the number of degrees that an object is panned or tilted when the cursor is dragged the entire width of the VR movie image). The default value is 180.0.
minPan	The minimum pan angle, in degrees. The value of this field must be less than the value of the <code>maxPan</code> field.

Media Data Atom Types

<code>maxPan</code>	The maximum pan angle, in degrees. The value of this field must be greater than the value of the <code>minPan</code> field.
<code>defaultPan</code>	The default pan angle, in degrees. This is the pan angle used when the object is first displayed. The value of this field must be greater than or equal to the value of the <code>minPan</code> field and less than or equal to the value of the <code>maxPan</code> field.
<code>minTilt</code>	The minimum tilt angle, in degrees. The default value is +90.0. The value of this field must be less than the value of the <code>maxTilt</code> field.
<code>maxTilt</code>	The maximum tilt angle, in degrees. The default value is -90.0. The value of this field must be greater than the value of the <code>minTilt</code> field.
<code>defaultTilt</code>	The default tilt angle, in degrees. This is the tilt angle used when the object is first displayed. The value of this field must be greater than or equal to the value of the <code>minTilt</code> field and less than or equal to the value of the <code>maxTilt</code> field.
<code>minFieldOfView</code>	The minimum field of view to which the object can zoom. The valid range for this field is from 1 to the value of the <code>fieldOfView</code> field. The value of this field must be positive.
<code>fieldOfView</code>	The image field of view, in degrees, for the entire object. The value in this field must be greater than or equal to the value of the <code>minFieldOfView</code> field.
<code>defaultFieldOfView</code>	The default field of view for the object. This is the field of view used when the object is first displayed. The value in this field must be greater than or equal to the value of the <code>minFieldOfView</code> field and less than or equal to the value of the <code>fieldOfView</code> field.
<code>defaultViewCenterH</code>	The default horizontal view center.
<code>defaultViewCenterV</code>	The default vertical view center.
<code>viewRate</code>	The view rate (that is, the positive or negative rate at which the view animation in the object plays, if view animation is enabled). The value of this field must be from -100.0 through +100.0, inclusive.

Media Data Atom Types

<code>frameRate</code>	The frame rate (that is, the positive or negative rate at which the frame animation in a view plays, if frame animation is enabled). The value of this field must be from -100.0 through $+100.0$, inclusive.
<code>animationSettings</code>	A set of 32-bit flags that encode information about the animation settings of the object.
<code>controlSettings</code>	A set of 32-bit flags that encode information about the control settings of the object.

Object Controller Types

The `movieType` field of the object sample atom structure specifies an object controller type, that is, the user interface to be used to manipulate the object.

QuickTime VR supports the following controller types:

```
enum ObjectUITypes {
    kGrabberScrollerUI          = 1,
    kOldJoyStickUI              = 2,
    kJoystickUI                  = 3,
    kGrabberUI                   = 4,
    kAbsoluteUI                  = 5
};
```

Constant descriptions

<code>kGrabberScrollerUI</code>	The default controller, which displays a hand for dragging and rotation arrows when the cursor is along the edges of the object window.
<code>kOldJoyStickUI</code>	A joystick controller, which displays a joystick-like interface for spinning the object. With this controller, the direction of panning is reversed from the direction of the grabber.
<code>kJoystickUI</code>	A joystick controller, which displays a joystick-like interface for spinning the object. With this controller, the direction of panning is consistent with the direction of the grabber.
<code>kGrabberUI</code>	A grabber-only interface, which displays a hand for dragging but does <i>not</i> display rotation arrows when the cursor is along the edges of the object window.

Media Data Atom Types

`kAbsoluteUI` An absolute controller, which displays a finger for pointing. The absolute controller switches views based on a row-and-column grid mapped into the object window.

Animation Settings

The `animationSettings` field of the object sample atom is a long integer that specifies a set of animation settings for an object node. Animation settings specify characteristics of the movie while it is playing. Use these constants to specify animation settings:

```
enum QTVRAnimationSettings {
    kQTVRObjectAnimateViewFramesOn          = (1 << 0),
    kQTVRObjectPalindromeViewFramesOn       = (1 << 1),
    kQTVRObjectStartFirstViewFrameOn        = (1 << 2),
    kQTVRObjectAnimateViewsOn               = (1 << 3),
    kQTVRObjectPalindromeViewsOn            = (1 << 4),
    kQTVRObjectSyncViewToFrameRate          = (1 << 5),
    kQTVRObjectDontLoopViewFramesOn         = (1 << 6),
    kQTVRObjectPlayEveryViewFrameOn        = (1 << 7)
};
```

Constant descriptions

`kQTVRObjectAnimateViewFramesOn`
If this bit is set, play all frames in the current view state.

`kQTVRObjectPalindromeViewFramesOn`
If this bit is set, play a back-and-forth animation of the frames of the current view state.

`kQTVRObjectStartFirstViewFrameOn`
If this bit is set, play the frame animation starting with the first frame in the view (that is, at the view start time).

`kQTVRObjectAnimateViewsOn`
If this bit is set, play all views of the current object in the default row of views.

`kQTVRObjectPalindromeViewsOn`
If this bit is set, play a back-and-forth animation of all views of the current object in the default row of views.

`kQTVRObjectSyncViewToFrameRate`
If this bit is set, synchronize the view animation to the

Media Data Atom Types

frame animation and use the same options as for frame animation.

`kQTVRObjectDontLoopViewFramesOn`

If this bit is set, stop playing the frame animation in the current view at the end.

`kQTVRObjectPlayEveryViewFrameOn`

If this bit is set, play every view frame regardless of play rate. The play rate is used to adjust the duration in which a frame appears but no frames are skipped so the rate is not exact.

Control Settings

The `controlSettings` field of the object sample atom is a long integer that specifies a set of control settings for an object node. Control settings specify whether the object can wrap during panning and tilting, as well as other features of the node. The control settings are specified using these bit flags:

```
enum QTVRControlSettings {
    kQTVRObjectWrapPanOn           = (1 << 0),
    kQTVRObjectWrapTiltOn         = (1 << 1),
    kQTVRObjectCanZoomOn          = (1 << 2),
    kQTVRObjectReverseHControlOn  = (1 << 3),
    kQTVRObjectReverseVControlOn  = (1 << 4),
    kQTVRObjectSwapHVControlOn    = (1 << 5),
    kQTVRObjectTranslationOn      = (1 << 6)
};
```

Constant descriptions

`kQTVRObjectWrapPanOn`

If this bit is set, enable wrapping during panning. When this control setting is enabled, the user can wrap around from the current pan constraint maximum value to the pan constraint minimum value (or vice versa) using the mouse or arrow keys.

`kQTVRObjectWrapTiltOn`

If this bit is set, enable wrapping during tilting. When this control setting is enabled, the user can wrap around from the current tilt constraint maximum value to the tilt

Media Data Atom Types

constraint minimum value (or vice versa) using the mouse or arrow keys.

`kQTVRObjectCanZoomOn`

If this bit is set, enable zooming. When this control setting is enabled, the user can change the current field of view using the zoom-in and zoom-out keys on the keyboard (or using the VR controller buttons).

`kQTVRObjectReverseHControlOn`

If this bit is set, reverse the direction of the horizontal control.

`kQTVRObjectReverseVControlOn`

If this bit is set, reverse the direction of the vertical control.

`kQTVRObjectSwapHVControlOn`

If this bit is set, exchange the horizontal and vertical controls.

`kQTVRObjectTranslationOn`

If this bit is set, enable translation. When this setting is enabled, the user can translate using the mouse when either the translate key is held down or the controller translation mode button is toggled on.

Track References for Object Tracks

The track references to an object's image and hot spot tracks are not handled the same way as track references to panoramas. The track reference types are the same (`kQTVRImageTrackRefType` and `kQTVRHotSpotTrackRefAtomType`), but the location of the reference indexes is different. There is no entry in the object sample atom for the track reference indexes. Instead, separate atoms using the `VRTrackRefEntry` structure are stored as siblings to the object sample atom. The types of these atoms are `kQTVRImageTrackRefAtomType` and `kQTVRHotSpotTrackRefAtomType`. If either of these atoms is not present, then the reference index to the corresponding track is assumed to be 1.

Note


The `trackResolution` field in the `VRTrackRefEntry` structure is ignored for object tracks. ♦

The actual views of an object for an object node are contained in an **object image track**, which is usually a standard QuickTime video track. (An object

image track can also be any type of track that is capable of displaying an image, such as a QuickTime 3D track.)

As described in Chapter 1 of *Programming With QuickTime VR*, these views are often captured by moving a camera around the object in a defined pattern of pan and tilt angles. The views must then be ordered into an object image array, which is stored as a one-dimensional sequence of frames in the movie's video track (see Figure 3-13).

Figure 3-13 The structure of an image track for an object



For object movies containing frame animation, each animated view in the object image array consists of the animating frames. It is not necessary that each view in the object image array contain the same number of frames, but the view duration of all views in the object movie must be the same.

For object movies containing alternate view states, alternate view states are stored as separate object image arrays that immediately follow the preceding view state in the object image track. Each state does not need to contain the same number of frames. However, the total movie time of each view state in an object node must be the same.

Movie Media

Movie media is used to encapsulate embedded movies within QuickTime movies. This feature is available in QuickTime 4.1.

Movie Sample Description

The movie media doesn't have a unique sample description. It uses the minimum sample description, which is `SampleDescriptionRecord`.

Movie Media Sample Format

Each sample in the movie media is a QuickTime atom container. All root level atoms and their contents are enumerated in the following list. Note that the contents of all atoms are stored in big-endian format.

`kMovieMediaDataReference`

This atom contains a data reference type and a data reference. The data reference type is stored as an `OSType` at the start of the atom. The data reference is stored following the data reference type. If the data reference type is URL and the data reference is for a movie on the Apple website, the contents of the atom would be url `http://www.apple.com/foo.mov`.

There may be more than one atom of this type. The first atom of this type should have an atom ID of 1. Additional data references should be numbered sequentially.

`kMovieMediaDefaultDataReferenceID`

This atom contains a `QTAtomID` that indicates the ID of the data reference to use when instantiating the embedded movie for this sample. If this atom is not present, the data reference with an ID of 1 is used.

`kMovieMediaSlaveTime`

This atom contains a Boolean that indicates whether or not the `TimeBase` of the embedded movie should be slaved to the `TimeBase` of the parent movie. If the `TimeBase` is slaved,

the embedded movie's zero time will correspond to the start time of its movie media sample. Further, the playback rate of the embedded movie will always be the same as the parent movie's. If the `TimeBase` is not slaved, the embedded movie will default to a rate of 0, and a default time of whatever default time value it instantiated with (which may not be 0). If the `TimeBase` is not slaved, the embedded movie can be played by either including an `AutoPlay` atom in the movie media sample or by using a wired action. If this atom is not present, the embedded movie defaults to not slaved.

`kMovieMediaSlaveAudio`

This atom contains a Boolean that indicates whether or not the audio properties of the embedded movie should be slaved to those of the parent movie. When audio is slaved, all audio properties of the containing track are duplicated in the embedded movie. These properties include sound volume, balance, bass and treble, and level metering. If this atom is not present, the embedded movie defaults to not slaved audio.

`kMovieMediaSlaveGraphicsMode`

This atom contains a Boolean that indicates how the graphics mode of the containing track is applied to the embedded movie. If the graphics mode is not slaved, then the entire embedded movie is imaged using its own graphics modes. The result of the drawing of the embedded movie is composited onto the containing movie using the graphics mode of the containing track. If the graphics mode is slaved, then the graphics mode of each track in the embedded movie is ignored and instead the graphics mode of the containing track is used. In this case, the tracks of the embedded movie composite their drawing directly into the parent movie's contents. If this atom is not present, the graphics mode defaults to not slaved. Graphics mode slaving is useful for compositing semi-transparent media—for example, a PNG with an alpha channel—on top of other media.

`kMovieMediaSlaveTrackDuration`

This atom contains a Boolean that indicates how the Movie Media Handler should react when the duration of the

embedded movie is different than the duration of the movie media sample that it is contained by. When the movie media sample is created, the duration of the embedded movie may not yet be known. Therefore, the duration of the media sample may not be correct. In this case, the Movie Media Handler can do one of two things. If this atom is not present or it contains a value of false, the Movie Media Handler will respect the duration of media sample that contains the embedded movie. If the embedded movie has a longer duration than the movie media sample, the embedded movie will be truncated to the duration of the containing movie media sample. If the embedded movie is shorter, there will be a gap after it is finished playing. If this atom contains a value of true, the duration of the movie media sample will be adjusted to match the actual duration of the embedded movie. Because it is not possible to change an existing media sample, this will cause a new media sample to be added to the movie and the track's edit list to be updated to reference the new sample instead of the original sample.

Note

When the duration of the embedded movie's sample is adjusted, by default no other tracks are adjusted. This can cause the overall temporal composition to change in unintended ways. To maintain the complete temporal composition, a higher-level data structure which describes the temporal relationships between the various tracks must also be included with the movie. ♦

`kMovieMediaAutoPlay`

This atom contains a Boolean that indicates whether or not the embedded movie should start playing immediately after being instantiated. This atom is only used if the `TimeBase` of the embedded movie is not slaved to the parent movie (see the `kMovieMediaSlaveTime` atom (page 246) for more information). If auto play is requested, the movie will be played at its preferred rate after being instantiated. If this atom is not present, the embedded movie will not automatically play.

`kMovieMediaLoop`

Media Data Atom Types

This atom contains a `UInt8` that indicates how the embedded movie should loop. This atom is only used if the `TimeBase` of the embedded movie is not slaved to the parent movie (see the `kMovieMediaSlaveTime` atom (page 246) for more information). If this atom contains a 0, or if this atom is not present, the embedded movie will not loop. If this atom contains a value of 1, the embedded movie loops normally—that is, when it reaches the end it loops back to the beginning. If this atom contains a value of 2, the embedded movie uses palindromic looping. All other values are reserved.

`kMovieMediaUseMIMEType`

This atom contains text (not a C string or a pascal string) that indicates the MIME type of the movie import component that should be used to instantiate this media. This is useful in cases where the data reference may not contain MIME type information. If this atom is not present, the MIME type of the data reference as determined at instantiation time is used. This atom is intended to allow content creators a method for working around MIME type binding problems. It should not typically be required, and should not be included in movie media samples by default.

`kMovieMediaTitle`

Currently unused. It would contain text indicating the name of the embedded movie.

`kMovieMediaAltText`

This atom contains text (not a C string or a pascal string) that is displayed to the user when the embedded movie is being instantiated or if the embedded movie cannot be instantiated. If this atom is not present, the name of the data reference (typically the file name) is used.

`kMovieMediaClipBegin`

This atom contains a `MovieMediaTimeRecord` that indicates the time of the embedded movie that should be used. The clip begin atom provides a way to specify that a portion of the beginning of the embedded movie should not be used. If this atom is not present, the beginning of the embedded movie is not changed. Note that this atom does not change the time at which the embedded movie begins playing in the parent movie's time line. If the time specified in the clip

Media Data Atom Types

begin atom is greater than the duration of the embedded movie, then the embedded movie will not play at all.

```
struct MovieMediaTimeRecord {
    wide           time;
    TimeScale      scale;
};
```

`kMovieMediaClipDuration`

This atom contains a `MovieMediaTimeRecord` that indicates the duration of the embedded movie that should be used. The clip duration atom is applied by removing media from end of the embedded movie. If the clip duration atom is not present, then no media is removed from the end of the embedded movie. In situations where the sample contains both a clip duration and a clip begin atom, the clip begin is applied first. If the clip duration specifies a value that is larger than the duration of the embedded movie, no change is made to the embedded movie.

`kMovieMediaEnableFrameStepping`

This atom contains a Boolean that indicates whether or not the embedded movie should be considered when performing step operations, specifically using the interesting time calls with the `nextTimeStep` flag. If this atom is not present or is set to `false`, the embedded movie is not included in step calculations. If the atom is set to `true`, it is included in step calculations.

`kMovieMediaBackgroundColor`

This atom contains an `RGBColor` that is used for filling the background when the movie is being instantiated or when it fails to instantiate.

`kMovieMediaRegionAtom`

This atom contains a number of child atoms, shown below, which describe how the Movie Media Handler should resize the embedded movie. If this atom is not present, the Movie Media Handler resizes the child movie to completely fill the containing track's box.

```
kMovieMediaSpatialAdjustment
```

This atom contains an `OStype` that indicates how the embedded movie should be scaled to fit the track box. If this atom is not present, the default value is `kMovieMediaFitFill`. These modes are all based on SMIL layout options.

`kMovieMediaFitClipIfNecessary`

If the media is larger than the track box, it will be clipped; if it is smaller, any additional area will be transparent.

`kMovieMediaFitFill`

The media will be scaled to completely fill the track box.

`kMovieMediaFitMeet`

The media is proportionally scaled so that it is entirely visible in the track box and fills the largest area possible without changing the aspect ratio.

`kMovieMediaFitSlice`

The media is scaled proportionally so that the smaller dimension is completely visible.

`kMovieMediaFitScroll`

Not currently implemented. It currently has the same behavior as

`kMovieMediaFitClipIfNecessary`. When implemented, it will have the behavior described in the SMIL specification for a scrolling layout element.

`kMovieMediaRectangleAtom`

This atom contains four child atoms that define a rectangle. Not all child atoms must be present: top and left must both appear together, width and height must both appear together. The dimensions contained in this rectangle are used in place of the track box when applying the contents of the spatial adjustment atom. If the top and left are not specified, the top and left of the containing track's box are

used. If the width and height are not specified, the width and height of the containing track's box are used. Each child atom contains a `UInt32`.

```
kMovieMediaTop  
kMovieMediaLeft  
kMovieMediaWidth  
kMovieMediaHeight
```

Basic Data Types

This chapter describes a number of common data types that are used in QuickTime files.

Language Code Values

Some elements of a QuickTime file may be associated with a particular spoken language. To indicate the language associated with a particular object, the QuickTime file format uses language codes from the Macintosh Script Manager. Table 4-1 lists some of the language codes supported by QuickTime.

Table 4-1 QuickTime language code values

Language	Value	Language	Value
English	0	Georgian	52
French	1	Moldavian	53
German	2	Kirghiz	54
Italian	3	Tajiki	55
Dutch	4	Turkmen	56
Swedish	5	Mongolian	57
Spanish	6	MongolianCyr	58
Danish	7	Pashto	59
Portuguese	8	Kurdish	60

Table 4-1 QuickTime language code values (continued)

Language	Value	Language	Value
Norwegian	9	Kashmiri	61
Hebrew	10	Sindhi	62
Japanese	11	Tibetan	63
Arabic	12	Nepali	64
Finnish	13	Sanskrit	65
Greek	14	Marathi	66
Icelandic	15	Bengali	67
Maltese	16	Assamese	68
Turkish	17	Gujarati	69
Croatian	18	Punjabi	70
Traditional Chinese	19	Oriya	71
Urdu	20	Malayalam	72
Hindi	21	Kannada	73
Thai	22	Tamil	74
Korean	23	Telugu	75
Lithuanian	24	Sinhalese	76
Polish	25	Burmese	77
Hungarian	26	Khmer	78
Estonian	27	Lao	79
Lettish	28	Vietnamese	80
Latvian	28	Indonesian	81
Saamisk	29	Tagalog	82
Lappish	29	MalayRoman	83
Faeroese	30	MalayArabic	84
Farsi	31	Amharic	85

Table 4-1 QuickTime language code values (continued)

Language	Value	Language	Value
		Tigrinya	86
Russian	32	Galla	87
Simplified Chinese	33	Oromo	87
Flemish	34	Somali	88
Irish	35	Swahili	89
Albanian	36	Ruanda	90
Romanian	37	Rundi	91
Czech	38	Chewa	92
Slovak	39	Malagasy	93
Slovenian	40	Esperanto	94
Yiddish	41	Welsh	128
Serbian	42	Basque	129
Macedonian	43	Catalan	130
Bulgarian	44	Latin	131
Ukrainian	45	Quechua	132
Byelorussian	46	Guarani	133
Uzbek	47	Aymara	134
Kazakh	48	Tatar	135
Azerbaijani	49	Uighur	136
AzerbaijanAr	50	Dzongkha	137
Armenian	51	JavaneseRom	138

Calendar Date and Time Values

QuickTime movies store date and time information in Macintosh date format: a 32-bit value indicating the number of seconds that have passed since midnight January 1, 1904.

Matrices

QuickTime files use matrices to describe spatial information about many objects, such as tracks within a movie.

A transformation matrix defines how to map points from one coordinate space into another coordinate space. By modifying the contents of a transformation matrix, you can perform several standard graphics display operations, including translation, rotation, and scaling. The matrix used to accomplish two-dimensional transformations is described mathematically by a 3-by-3 matrix.

All values in the matrix are 32-bit fixed-point numbers divided as 16.16, except for the {u, v, w} column, which contains 32-bit fixed-point numbers divided as 2.30. Figure 4-1 depicts how QuickTime uses matrices to transform displayed objects.

Figure 4-1 How display matrices are used in QuickTime

$$\begin{bmatrix} x & y & 1 \end{bmatrix} \times \begin{bmatrix} a & b & u \\ c & d & v \\ t_x & t_y & w \end{bmatrix} = \begin{bmatrix} x' & y' & 1 \end{bmatrix}$$

Figure 4-2 Applying the transform

$$\begin{array}{l}
 p = ax + cy + tx \\
 q = bx + dy + ty \\
 r = ux + vy + w
 \end{array}
 \quad
 \begin{bmatrix} x' & y' & 1 \end{bmatrix}
 =
 \frac{\begin{bmatrix} p & q & r \end{bmatrix}}{r}$$

Graphics Modes

QuickTime files use graphics modes to describe how one video or graphics layer should be combined with the layers beneath it. Graphics modes are also known as transfer modes. Some graphics modes require a color to be specified for certain operations, such as blending to determine the blend level. QuickTime uses the graphics modes defined by Apple's QuickDraw.

The most common graphics modes are `noErase` and `ditherCopy`, which simply indicate that the image should not blend with the image behind it, but overwrite it. QuickTime also defines several additional graphics modes.

Table 4-2 lists the additional graphics modes supported by QuickTime.

Table 4-2 QuickTime graphics modes

Mode	Uses opcolor	Code	Description
Copy		0x0	Copy the source image over the destination.
Dither copy		0x40	Dither the image (if needed), otherwise do a copy.
Blend	yes	0x20	Replaces destination pixel with a blend of the source and destination pixel colors, with the proportion for each channel controlled by that channel in the opcolor.
Transparent	yes	0x24	Replaces the destination pixel with the source pixel if the source pixel isn't equal to the opcolor.
Straight alpha		0x100	Replaces the destination pixel with a blend of the source and destination pixels, with the proportion controlled by the alpha channel.
Premul white alpha		0x101	Pre-multiplied with white means that the color components of each pixel have already been blended with a white pixel, based on their alpha channel value. Effectively, this means that the image has already been combined with a white background. First, remove the white from each pixel and then blend the image with the actual background pixels.

Table 4-2 QuickTime graphics modes

Mode	Uses opcolor	Code	Description
Premul black alpha		0x102	Pre-multiplied with black is the same as pre-multiplied with white, except the background color that the image has been blended with is black instead of white.
Straight alpha blend	yes	0x104	Similar to straight alpha, but the alpha value used for each channel is the combination of the alpha channel and that channel in the opcolor.
Composition (dither copy)		0x103	(Tracks only) The track is drawn offscreen, and then composed onto the screen using dither copy

RGB Colors

Many atoms in the QuickTime file format contain RGB color values. These are usually stored as three consecutive unsigned 16-bit integers in the following order: red, green, blue.

Balance

Balance values are represented as 16-bit, fixed-point numbers that range from -1.0 to +1.0. The high-order 8 bits contain the integer portion of the value; the low-order 8 bits contain the fractional part. Negative values weight the balance toward the left speaker; positive values emphasize the right channel. Setting the balance to 0 corresponds to a neutral setting.

CHAPTER 4

Basic Data Types

Some Useful Examples and Scenarios

This chapter contains a number of examples that can help you pull together all of the material in this book by examining the atom structure that results from a number of different scenarios.

The chapter is divided into the following topics:

- “Creating, Copying, and Disposing of Atom Containers” (page 262) discusses the various ways you can work with atom containers, along with illustrations and sample code that show usage.
- “Creating an Effect Description” (page 272) discusses how you create an effect description by creating an atom container, inserting a QT atom that specifies the effect, and inserting a set of QT atoms that set its parameters.
- “Creating Movies With Modifier Tracks” (page 280) provides sample code showing you how to create a movie with modifier tracks.
- “Authoring Movies With External Movie Targets” (page 281) discusses how to author movies with external targets, using two new target atoms introduced in QuickTime 4.
- “Adding Wired Actions To a Flash Track” (page 284) explains the steps you need to follow in order to add wired actions to a Macromedia Flash track.
- “Creating Video Tracks at 30 Frames-per-Second” (page 287) discusses creating 30 fps video.
- “Creating Video Tracks at 29.97 Frames-per-Second” (page 287) describes creating 29.97 fps video.
- “Creating Audio Tracks at 44.1 Khz” (page 288) provides an example of creating an audio track.
- “Creating a Timecode Track for 29.97 FPS Video” (page 289) presents a timecode track example.
- “Playing With Edit Lists” (page 292) discusses how to interpret edit list data.

- “Interleaving Movie Data” (page 294) shows how a movie’s tracks are interleaved in the movie data file.
- “Referencing Two Data Files With a Single Track” (page 296) shows how track data may reside in more than one file.
- “Getting the Name of a QuickTime VR Node” (page 298) discusses how you can use standard QuickTime atom container functions to retrieve the information in a QuickTime VR node header atom.
- “Adding Custom Atoms in a QuickTime VR Movie” (page 299) describes how to add custom atoms to either the QuickTime VR world or node information atom containers.
- “Adding Atom Containers in a QuickTime VR Movie” (page 301) shows the code you would use to add VR world and node information atom containers to a QTVR track.
- “Optimizing QuickTime VR Movies for Web Playback” (page 302) describes how to use the QTVR Flattener, a movie export component that converts an existing QuickTime VR single node movie into a new movie that is optimized for viewing on the Web.

Creating, Copying, and Disposing of Atom Containers

Before you can add atoms to an atom container, you must first create the container by calling `QTNewAtomContainer`. The code sample shown in Listing 5-1 calls `QTNewAtomContainer` to create an atom container.

Listing 5-1 Creating a new atom container

```
QTAtomContainer spriteData;  
OSErr err  
// create an atom container to hold a sprite's data  
err=QTNewAtomContainer (&spriteData);
```

When you have finished using an atom container, you should dispose of it by calling the `QTDisposeAtomContainer` function. The sample code shown in

Some Useful Examples and Scenarios

Listing 5-2 calls `QTDisposeAtomContainer` to dispose of the `spriteData` atom container.

Listing 5-2 Disposing of an atom container

```
if (spriteData)
    QTDisposeAtomContainer (spriteData);
```

Creating New Atoms


You can use the `QTInsertChild` function to create new atoms and insert them in a QT atom container. The `QTInsertChild` function creates a new child atom for a parent atom. The caller specifies an atom type and atom ID for the new atom. If you specify a value of 0 for the atom ID, `QTInsertChild` assigns a unique ID to the atom.

`QTInsertChild` inserts the atom in the parent's child list at the index specified by the `index` parameter; any existing atoms at the same index or greater are moved toward the end of the child list. If you specify a value of 0 for the `index` parameter, `QTInsertChild` inserts the atom at the end of the child list.

The code sample in Listing 5-3 creates a new QT atom container and calls `QTInsertChild` to add an atom. The resulting QT atom container is shown in Figure 5-1. The offset value 10 is returned in the `firstAtom` parameter.

Listing 5-3 Creating a new QT atom container and calling `QTInsertChild` to add an atom.

```
QTAtom firstAtom;
QTAtomContainer container;
OSErr err
err = QTNewAtomContainer (&container);
if (!err)
    err = QTInsertChild (container, kParentAtomIsContainer, 'abcd',
        1000, 1, 0, nil, &firstAtom);
```


Figure 5-1 QT atom container after inserting an atom

The following code sample calls `QTInsertChild` to create a second child atom. Because a value of 1 is specified for the `index` parameter, the second atom is inserted in front of the first atom in the child list; the index of the first atom is changed to 2. The resulting QT atom container is shown in Figure 5-2.

```

QTAtom secondAtom;

FailOSErr (QTInsertChild (container, kParentAtomIsContainer, 'abcd',
    2000, 1, 0, nil, &secondAtom));
  
```

Figure 5-2 QT atom container after inserting a second atom

Some Useful Examples and Scenarios

You can call the `QTFindChildByID` function to retrieve the changed offset of the first atom that was inserted, as shown in the following example. In this example, the `QTFindChildByID` function returns an offset of 20.

```
firstAtom = QTFindChildByID (container, kParentAtomIsContainer, 'abcd',
                             1000, nil);
```

Listing 5-4 shows how the `QTInsertChild` function inserts a leaf atom into the atom container `sprite`. The new leaf atom contains a sprite image index as its data.


Listing 5-4 Inserting a child atom

```
if ((propertyAtom = QTFindChildByIndex (sprite, kParentAtomIsContainer,
                                         kSpritePropertyImageIndex, 1, nil)) == 0)

    FailOSErr (QTInsertChild (sprite, kParentAtomIsContainer,
                             kSpritePropertyImageIndex, 1, 1, sizeof(short), &imageIndex,
                             nil));
```

Copying Existing Atoms

QuickTime provides several functions for copying existing atoms within an atom container. The `QTInsertChildren` function inserts a container of atoms as children of a parent atom in another atom container. Figure 5-3 shows two example QT atom containers, A and B.

Figure 5-3 Two QT atom containers, A and B


The following code sample calls `QTFindChildByID` to retrieve the offset of the atom in container A. Then, the code sample calls the `QTInsertChildren` function to insert the atoms in container B as children of the atom in container A. Figure 5-4 shows what container A looks like after the atoms from container B have been inserted.

```

QTAtom targetAtom;

targetAtom = QTFindChildByID (containerA, kParentAtomIsContainer, 'abcd',
    1000, nil);

FailOSErr (QTInsertChildren (containerA, targetAtom, containerB));
  
```

Figure 5-4 QT atom container after child atoms have been inserted

In Listing 5-5, the `QTInsertChild` function inserts a parent atom into the atom container `theSample`. Then, the code calls `QTInsertChildren` to insert the container `theSprite` into the container `theSample`. The parent atom is `newSpriteAtom`.

Listing 5-5 Inserting a container into another container

```

FailOSErr (QTInsertChild (theSample, kParentAtomIsContainer,
    kSpriteAtomType, spriteID, 0, 0, nil, &newSpriteAtom));

FailOSErr (QTInsertChildren (theSample, newSpriteAtom, theSprite));

```

QuickTime provides three other functions you can use to manipulate atoms in an atom container. The `QTReplaceAtom` function replaces an atom and its children with a different atom and its children. You can call the `QTSwapAtoms` function to swap the contents of two atoms in an atom container; after swapping, the ID and index of each atom remains the same. The `QTCopyAtom` function copies an atom and its children to a new atom container.

Retrieving Atoms From an Atom Container

QuickTime provides functions you can use to retrieve information about the types of a parent atom's children, to search for a specific atom, and to retrieve a leaf atom's data.

You can use the `QTCountChildrenOfType` and `QTGetNextChildType` functions to retrieve information about the types of an atom's children. The `QTCountChildrenOfType` function returns the number of children of a given atom type for a parent atom. The `QTGetNextChildType` function returns the next atom type in the child list of a parent atom.

You can use the `QTFindChildByIndex`, `QTFindChildByID`, and `QTNextChildAnyType` functions to retrieve an atom. You call the `QTFindChildByIndex` function to search for and retrieve a parent atom's child by its type and index within that type.

Listing 5-6 shows the sample code function `SetSpriteData`, which updates an atom container that describes a sprite. (For more information about sprites and the Sprite Toolbox, refer to the book *Programming With Wired Movies and Sprite Animation*, available at <http://developer.apple.com/techpubs/quicktime/qtdevdocs/RM/PDF.htm>.) For each property of the sprite that needs to be updated, `SetSpriteData` calls `QTFindChildByIndex` to retrieve the appropriate atom from the atom container. If the atom is found, `SetSpriteData` calls `QTSetAtomData` to replace the atom's data with the new value of the property. If the atom is not found, `SetSpriteData` calls `QTInsertChild` to add a new atom for the property.

Listing 5-6 Finding a child atom by index

```
OSErr SetSpriteData (QTAtomContainer sprite, Point *location,
    short *visible, short *layer, short *imageIndex)
{
    OSErr err = noErr;
    QTAtom propertyAtom;

    // if the sprite's visible property has a new value
    if (visible)
    {
        // retrieve the atom for the visible property --
        // if none exists, insert one
        if ((propertyAtom = QTFindChildByIndex (sprite,
            kParentAtomIsContainer, kSpritePropertyVisible, 1,
```

Some Useful Examples and Scenarios

```

        nil)) == 0)
        FailOSErr (QTInsertChild (sprite, kParentAtomIsContainer,
                                kSpritePropertyVisible, 1, 1, sizeof(short), visible,
                                nil))

        // if an atom does exist, update its data
    else
        FailOSErr (QTSetAtomData (sprite, propertyAtom,
                                sizeof(short), visible));
    }

    // ...
    // handle other sprite properties
    // ...
}

```

You can call the `QTFindChildByID` function to search for and retrieve a parent atom's child by its type and ID. The sample code function `AddSpriteToSample`, shown in Listing 5-7, adds a sprite, represented by an atom container, to a key sample, represented by another atom container. `AddSpriteToSample` calls `QTFindChildByID` to determine whether the atom container `theSample` contains an atom of type `kSpriteAtomType` with the ID `spriteID`. If not, `AddSpriteToSample` calls `QTInsertChild` to insert an atom with that type and ID. A value of 0 is passed for the `index` parameter to indicate that the atom should be inserted at the end of the child list. A value of 0 is passed for the `dataSize` parameter to indicate that the atom does not have any data. Then, `AddSpriteToSample` calls `QTInsertChildren` to insert the atoms in the container `theSprite` as children of the new atom. `FailIf` and `FailOSErr` are macros that exit the current function when an error occurs.

Listing 5-7 Finding a child atom by ID

```

OSErr AddSpriteToSample (QTAtomContainer theSample,
                        QTAtomContainer theSprite, short spriteID)
{
    OSErr err = noErr;
    QTAtom newSpriteAtom;

    FailIf (QTFindChildByID (theSample, kParentAtomIsContainer,
                            kSpriteAtomType, spriteID, nil), paramErr);
}

```

Some Useful Examples and Scenarios

```

    FailOSErr (QTInsertChild (theSample, kParentAtomIsContainer,
        kSpriteAtomType, spriteID, 0, 0, nil, &newSpriteAtom));
    FailOSErr (QTInsertChildren (theSample, newSpriteAtom, theSprite));
}

```

Once you have retrieved a child atom, you can call `QTNextChildAnyType` function to retrieve subsequent children of a parent atom. `QTNextChildAnyType` returns an offset to the next atom of any type in a parent atom's child list. This function is useful for iterating through a parent atom's children quickly.

QuickTime also provides functions for retrieving an atom's type, ID, and data. You can call `QTGetAtomTypeAndID` function to retrieve an atom's type and ID. You can access an atom's data in one of three ways.

- To copy an atom's data to a handle, you can use the `QTCopyAtomDataToHandle` function.
- To copy an atom's data to a pointer, you can use the `QTCopyAtomDataToPtr` function.
- To access an atom's data directly, you should lock the atom container in memory by calling `QTLockContainer`. Once the container is locked, you can call `QTGetAtomDataPtr` to retrieve a pointer to an atom's data. When you have finished accessing the atom's data, you should call the `QTUnlockContainer` function to unlock the container in memory.

Modifying Atoms

QuickTime provides functions that you can call to modify attributes or data associated with an atom in an atom container. To modify an atom's ID, you call the function `QTSetAtomID`.

You use the `QTSetAtomData` function to update the data associated with a leaf atom in an atom container. The `QTSetAtomData` function replaces a leaf atom's data with new data. The code sample in Listing 5-8 calls

`QTFindChildByIndex` to determine whether an atom container contains a sprite's visible property. If so, the sample calls `QTSetAtomData` to replace the atom's data with a new visible property.

Listing 5-8 Modifying an atom's data

```

QTAtom propertyAtom;

// if the atom isn't in the container, add it
if ((propertyAtom = QTFindChildByIndex (sprite, kParentAtomIsContainer,
    kSpritePropertyVisible, 1, nil)) == 0)
    Fail0SErr (QTInsertChild (sprite, kParentAtomIsContainer,
        kSpritePropertyVisible, 1, 0, sizeof(short), visible, nil))

// if the atom is in the container, replace its data
else
    Fail0SErr (QTSetAtomData (sprite, propertyAtom, sizeof(short),
        visible));

```

Removing Atoms From an Atom Container

To remove atoms from an atom container, you can use the `QTRemoveAtom` and `QTRemoveChildren` functions. The `QTRemoveAtom` function removes an atom and its children, if any, from a container. The `QTRemoveChildren` function removes an atom's children from a container, but does not remove the atom itself. You can also use `QTRemoveChildren` to remove all the atoms in an atom container. To do so, you should pass the constant `kParentAtomIsContainer` for the atom parameter.

The code sample shown in Listing 5-9 adds override samples to a sprite track to animate the sprites in the sprite track. The `sample` and `spriteData` variables are atom containers. The `spriteData` atom container contains atoms that describe a single sprite. The `sample` atom container contains atoms that describes an override sample.

Each iteration of the `for` loop calls `QTRemoveChildren` to remove all atoms from both the `sample` and the `spriteData` containers. The sample code updates the index of the image to be used for the sprite and the sprite's location and calls `SetSpriteData` (Listing 5-6), which adds the appropriate atoms to the `spriteData` atom container. Then, the sample code calls `AddSpriteToSample` (Listing 5-7) to add the `spriteData` atom container to the `sample` atom container. Finally, when all the sprites have been updated, the sample code calls `AddSpriteSampleToMedia` to add the override sample to the sprite track.

Listing 5-9 Removing atoms from a container

```

QAtomContainer sample, spriteData;

// ...
// add the sprite key sample
// ...

// add override samples to make the sprites spin and move
for (i = 1; i <= kNumOverrideSamples; i++)
{
    QTRemoveChildren (sample, kParentAtomIsContainer);
    QTRemoveChildren (spriteData, kParentAtomIsContainer);

    // ...
    // update the sprite:
    // - update the imageIndex
    // - update the location
    // ...

    // add atoms to spriteData atom container
    SetSpriteData (spriteData, &location, nil, nil, &imageIndex);

    // add the spriteData atom container to sample
    err = AddSpriteToSample (sample, spriteData, 2);

    // ...
    // update other sprites
    // ...

    // add the sample to the media
    err = AddSpriteSampleToMedia (newMedia, sample,
    kSpriteMediaFrameDuration, false);
}

```

Creating an Effect Description

An effect description tells QuickTime which effect to execute and contains the parameters that control how the effect behaves at runtime. You create an effect description by creating an atom container, inserting a QT atom that specifies the effect, and inserting a set of QT atoms that set its parameters.

There are support functions you can call to assist you in this process.

`QTCreateStandardParameterDialog` returns a complete effect description that you

can use, including user-selected settings; you only need to add `kEffectSourceName` atoms to the description for effects that require sources. At a lower level, `QTGetEffectsList` returns a list of the available effects and `ImageCodecGetParameterList` will return a description of the parameters for an effect, including the default value for each parameter in the form of a QT atom that can be inserted directly into an effect description.

Structure of an Effect Description

An effect description is the sole media sample for an effect track. An effect description is implemented as a `QTAtomContainer` structure, the general QuickTime structure for holding a set of QuickTime atoms. All effect descriptions must contain the set of required atoms, which specify attributes such as which effect component to use. In addition, effect descriptions can contain a variable number of parameter atoms, which hold the values of the parameters for the effect.

Each atom contains either data or a set of child atoms. If a parameter atom contains data, the data is the value of the parameter, and this value remains constant while the effect executes. If a parameter atom contains a set of child atoms, they typically contain a tween entry so the value of the parameter will be interpolated for the duration of the effect.

You assemble an effect description by adding the appropriate set of atoms to a `QTAtomContainer` structure.

You can find out what the appropriate atoms are by making an `ImageCodecGetParameterList` call to the effect component. This fills an atom container with a set of parameter description atoms. These atoms contain descriptions of the effect parameters, such as each parameter's atom type, data range, default value, and so on. The default value in each description atom is itself a `QTAtom` that can be inserted directly into your effect description.

You can modify the data in the parameter atoms directly, or let the user set them by calling `QTCreateStandardParameterDialog`, which returns a complete effect description (you need to add `kEffectSourceName` atoms for effects that require sources).

You then add the effect description to the media of the effect track.

Required Atoms of an Effects Description

There are several required atoms that an effect description must contain. The first is the `kParameterWhatName` atom. The `kParameterWhatName` atom contains the name of the effect. This specifies which of the available effects to use.

The code snippet shown in Listing 5-10 adds a `kParameterWhatName` atom to the atom container `effectDescription`. The constant `kCrossFadeTransitionType` contains the name of the cross-fade effect.

Listing 5-10 Adding a `kParameterWhatName` atom to the atom container `effectDescription`

```
effectCode = kCrossFadeTransitionType;
QTInsertChild(effectDescription, kParentAtomIsContainer,
              kParameterWhatName, kParameterWhatID, 0,
              sizeof(effectCode), &effectCode, nil);
```

In addition to the `kParameterWhatName` atom, the effect description for an effect that uses sources must contain one or more `kEffectSourceName` atoms. Each of these atoms contains the name of one of the effect's sources. An input map is used to map these names to the actual tracks of the movie that are the sources. "Creating an Input Map" (page 277) describes how to create the input map.

Parameter Atoms of an Effects Description

In addition to the required atoms, the effects description contains a variable number of parameter atoms. The number and types of parameter atoms vary from effect to effect. For example, the cross fade effect has only one parameter, while the general convolution filter effect has nine. Some effects have no parameters at all, and do not require any parameter atoms.

You can obtain the list of parameter atoms for a given effect by calling the effect component using `ImageCodecGetParameterList`. The parameter description atoms it returns include default settings for each parameter in the form of parameter atoms that you can insert into your effect description.

The `QTInsertChild` function is used to add these parameters to the effect description, as seen in the code example in Listing 5-10.

Some Useful Examples and Scenarios

Consider, for instance, the push effect. Its effect description contains a `kParameterWhatName` atom, two `kEffectSourceName` atoms, and two parameter atoms, one of which is a tween.

The `kParameterWhatName` atom specifies that this is a 'push' effect.

The two `kEffectSourceName` atoms specify the two sources that this effect will use, in this case 'srcA' and 'srcB'. The names correspond to entries in the effect track's input map.

The 'pcnt' parameter atom defines which frames of the effect are shown. This parameter contains a tween entry, so that the value of this parameter is interpolated as the effect runs. The interpolation of the 'pcnt' parameter causes consecutive frames of the effect to be rendered, creating the push effect.

The 'from' parameter determines the direction of the push. This parameter is set from an enumeration list, with 2 being defined as the bottom of the screen.

In this example, the source 'srcB' will push in from the bottom, covering the source 'srcA'.

The 'pcnt' parameter is normally tweened from 0 to 100, so that the effect renders completely, from 0 to 100 percent. In this example, the 'pcnt' parameter is tweened from 25 to 75, so the effect will start 25% of the way through (with 'srcB' already partly on screen) and finish 75% of the way through (with part of 'srcA' still visible).

Figure 5-5 shows the set of atoms that must be added to the entry description.

Figure 5-5 An example effect description for the Push effect

	Required atoms	Byte			
Use the effect component with the name 'push'	<table><tr><td>kParameterWhatName</td></tr><tr><td>'push'</td></tr></table>	kParameterWhatName	'push'	1	
kParameterWhatName					
'push'					
The first source is 'srcA' which is the name of a source defined in the input map	<table><tr><td>kEffectSourceName</td></tr><tr><td>'srcA'</td></tr></table>	kEffectSourceName	'srcA'	1	
kEffectSourceName					
'srcA'					
The second source is 'srcB' from the input map.	<table><tr><td>kEffectSourceName</td></tr><tr><td>'srcB'</td></tr></table>	kEffectSourceName	'srcB'	2	
kEffectSourceName					
'srcB'					
Parameter atoms					
The percentage value, which is tweened for 25% to 75%	<table><tr><td>'pcnt'</td></tr></table>	'pcnt'	1		
'pcnt'					
	<table><tr><td>kTweenType</td></tr><tr><td>kParamTypeDataFixed</td></tr></table>	kTweenType	kParamTypeDataFixed	1	
kTweenType					
kParamTypeDataFixed					
	<table><tr><td>kTweenData</td></tr><tr><td>0.25</td></tr><tr><td>0.75</td></tr></table>	kTweenData	0.25	0.75	1
kTweenData					
0.25					
0.75					
The direction from which the second source pushes the first. The value 2 indicates it pushes in from below.	<table><tr><td>'from'</td></tr><tr><td>2</td></tr></table>	'from'	2	1	
'from'					
2					

An important property of effect parameters is that most can be tweened (and some must be tweened). Tweening is QuickTime's general purpose interpolation mechanism (see "Tween Media" (page 168) for more information). For many parameters, it is desirable to allow the value of the parameter to change as the effect executes. In the example shown in Figure 5-5, the 'pcnt' parameter must be a tween. This parameter controls which frame of the effect is

rendered at any given time, so it must change for the effect to progress. The 'from' parameter is not a tween in the example above, but it could be if we wanted the direction of the push to change during the course of the effect.

Creating an Input Map

The input map is another `QTAtomContainer` that you attach to the effects track. It describes the sources used in the effect and gives a name to each source. This name is used to refer to the source in the effect description.

An input map works in concert with track reference atoms in the source tracks. A track reference atom of type `kTrackModifierReference` is added to each source track, which causes that source track's output to be redirected to the effects track. An input map is added to the effects track to identify the source tracks and give a name to each source, such as 'srcA' and 'srcB'. The effect can then refer to the sources by name, specifying that 'srcB' should slide in over 'srcA', for example.

Structure of an Input Map

The input map contains a set of atoms that refer to the tracks used as sources for the effect. Each source track is represented by one track reference atom of type `kTrackModifierInput`.

Each modifier input atom contains two children, one of type `kEffectDataSourceType`, and one of type `kTrackModifierType`, which hold the name and type of the source.

The name of the source is a unique identifier that you create, which is used in the effect description to reference the track. Any four-character name is valid, as long as it is unique in the set of source names.

IMPORTANT

Apple recommends you adopt the standard naming convention 'srcX', where X is a letter of the alphabet. Thus, your first source would be named 'srcA', the second 'srcB', and so forth. This convention is used here in this chapter. ▲

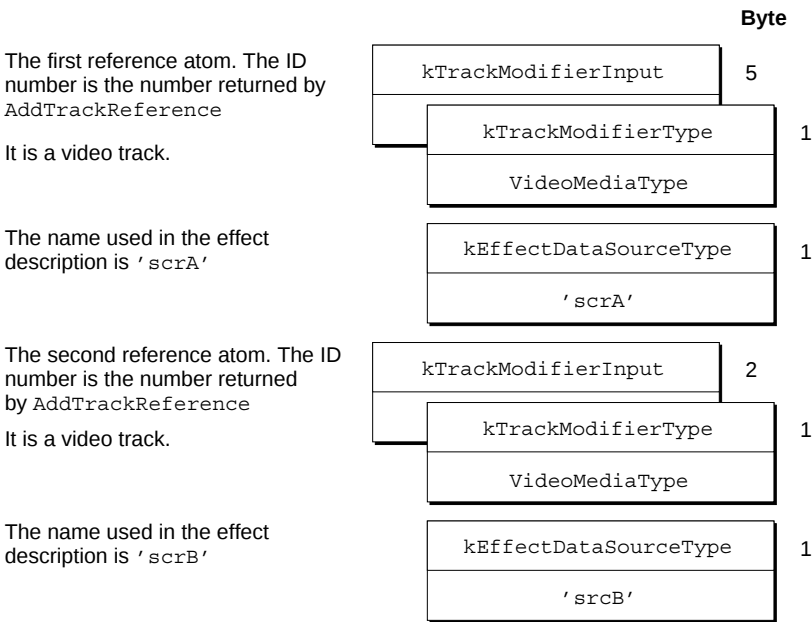
The child atom of type `kTrackModifierType` indicates the type of the track being referenced. For a video track the type is `VideoMediaType`, for a sprite track it is `SpriteMediaType`, and so forth. Video tracks are the most common track type

used as sources for effects. Only tracks that have a visible output, such as video and sprite tracks, can be used as sources for an effect. This means, for example, that sound tracks cannot be sources for an effect.

Figure 5-6 shows a completed input map that references two sources. The first source is a video track and is called 'srcA'. The second source, also a video track, is called 'srcB'.

You refer to a `kTrackModifierInput` atom by its index number, which is returned by the `AddTrackReference` function when you create the atom.

Figure 5-6 An example input map referencing two sources



Building Input Maps

The first step in creating an input map is to create a new `QTAtomContainer` to hold the map. You use the standard QuickTime container creation function:

```
QTNewAtomContainer(&inputMap);
```

Some Useful Examples and Scenarios

For each source you are creating, you need to call the `AddTrackReference` function. The track IDs of the effects track and the source track are passed as parameters to `AddTrackReference`, which creates an atom of type `kTrackModifierReference` and returns an index number. You use this index as the ID of the atom when you need to refer to it. You then insert the reference into the input map as an atom of type `kTrackModifierInput`.

The code in Listing 5-11 creates a reference to the track `firstSourceTrack`, and adds it to the input map.

Listing 5-11 Adding an input reference atom to an input map

```
AddTrackReference(theEffectsTrack, firstSourceTrack,
                  kTrackModifierReference, &referenceIndex);

QTInsertChild(inputMap, kParentAtomIsContainer,
              kTrackModifierInput, referenceIndex, 0, 0, nil, &inputAtom);
```

The `QTInsertChild` function returns the offset of the new modifier input atom in the `inputAtom` parameter.

You now need to add the name and type of the source track to the modifier input atom. Again, calling the `QTInsertChild` function does this, as shown in the following code snippet:

```
inputType = VideoMediaType;
QTInsertChild(inputMap, inputAtom,
              kTrackModifierType, 1, 0, sizeof(inputType), &inputType,
              nil);

aType = 'srcA';
QTInsertChild(inputMap, inputAtom, kEffectDataSourceType, 1, 0,
              sizeof(aType), &aType, nil);
```

This process is repeated for each source for the effect.

Creating Movies With Modifier Tracks

QuickTime 2.1 added additional functionality for media handlers. By way of modifier tracks, a media handler can send its data to another media handler rather than presenting its media directly. See “Modifier Tracks” (page 180) for a complete discussion of this feature.

To create a movie with modifier tracks, first you create a movie with all the desired tracks, then you create the modifier track. To link the modifier track to the track that it modifies, you use the `AddTrackReference` function as shown in Listing 5-12.

Listing 5-12 Linking a modifier track to the track it modifies

```
long addedIndex;  
AddTrackReference(aVideoTrack, aModifierTrack,  
                  kTrackModifierReference, &addedIndex);
```

The reference doesn’t completely describe the modifier track’s relationship to the track it modifies. Instead, the reference simply tells the modifier track to send its data to the specified track. The receiving track doesn’t “know” what it should do with that data. A single track may also be receiving data from more than one modifier track.

To describe how each modifier input should be used, each track’s media also has an input map. The media’s input map describes how the data being sent to each input of a track should be interpreted by the receiving track. After creating the reference, it is necessary to update the receiving track’s media input map. When `AddTrackReference` is called, it returns the index of the reference added. That index is the index of the input that needs to be described in the media input map. If the modifier track created above contains regions to change the shape of the video track, the code shown in Listing 5-13 updates the input map appropriately.

Listing 5-13 Updating the input map

```

QTAtomContainer inputMap;
QTAtom inputAtom;
OSType inputType;

Media aVideoMedia = GetTrackMedia(aVideoTrack);
GetMediaInputMap (aVideoMedia, &inputMap);

QTInsertChild(inputMap, kParentAtomIsContainer, kTrackModifierInput,
              addedIndex, 0,0, nil, &inputAtom);

inputType = kTrackModifierTypeClip;
QTInsertChild (inputMap, inputAtom, kTrackModifierType, 1, 0,
              sizeof(inputType), &inputType, nil);

SetMediaInputMap(aVideoMedia, inputMap);
QTDisposeAtomContainer(inputMap);

```

The media input map allows you to store additional information for each input. In the preceding example, only the type of the input is specified. In other types of references, you may need to specify additional data.

When a modifier track is playing an empty track edit, or is disabled or deleted, all receiving tracks are notified that the track input is inactive. When an input becomes inactive, it is reset to its default value. For example, if a track is receiving data from a clip modifier track and that input becomes inactive, the shape of the track reverts to the shape it would have if there were no clip modifier track.

Authoring Movies With External Movie Targets

QuickTime 4 enables you to author movies with external movie targets. To specify an action that targets an element of an external movie, you must identify the external movie by either its name or its ID. Two new target atom types have been introduced for this purpose; these atoms are used in addition to the existing target atoms, which you may use to specify that the element is a particular track or object within a track, such as a sprite.

Some Useful Examples and Scenarios

Note

A movie ID may be specified by an expression. ♦

The additional target atoms provided in QuickTime 4:

```
[(ActionTargetAtoms)] =
    <kActionTarget>

        <kTargetMovieName>
            [Pstring MovieName]
        OR
        <kTargetMovieID>
            [long MovieID]
        OR
        [(kExpressionAtoms)]
```

To tag a movie with a name or ID, you add a user data item of type 'plug' to the movie's user data. The index of the user data does not matter. The data specifies the name or ID.

You add a user data item of type 'plug' to the movie's user data with its data set to

```
"Movieid=MovieName"
```

where `MovieName` is the name of the movie.

You add a user data item of type 'plug' to the movie's user data with its data set to

```
"Movieid=MovieID"
```

where the ID is a signed long integer.

The QuickTime plug-in additionally supports EMBED tag parameters, which allow you to override a movie's name or ID within an HTML page.

Target Atoms for Embedded Movies

QuickTime 4.1 introduced target atoms to accommodate the addition of embedded movies. These target atoms allow for paths to be specified in a hierarchical movie tree.

Some Useful Examples and Scenarios

Target movies may be an external movie, the default movie, or any movie embedded within another movie. Targets are specified by using a movie path that may include parent and child movie relationships, and may additionally include track and track object target atoms as needed.

By using embedded `kActionTarget` atoms along with parent and child movie target atoms, you can build up paths for movie targets. Note that QuickTime looks for these embedded `kActionTarget` atoms only when evaluating a movie target, and any movie target type may contain a sibling `kActionTarget` atom.

Paths begin from the current movie, which is the movie containing the object that is handling an event. You may go up the tree using a `kTargetParentMovie` atom, or down the tree using one of five new child movie atoms. You may use a `kTargetRootMovie` atom as a shortcut to get to the top of the tree containing an embedded movie and may use the `movieByName` and `movieByID` atoms to specify a root external movie.

The target atoms are

- `kTargetRootMovie` (leaf atom, no data). This is the root movie containing the action handler.
- `kTargetParentMovie` (leaf atom, no data). This is the parent movie.

Note that there are five ways to specify an embedded child movie. Three of them specify movie track properties. Two specify properties of the currently loaded movie in a movie track.

- `kTargetChildMovieTrackName`. A child movie track specified by track name.
- `kTargetChildMovieTrackID`. A child movie track specified by track ID.
- `kTargetChildMovieTrackIndex`. A child movie track specified by track index.
- `kTargetChildMovieMovieName`. A child movie specified by the currently loaded movie's movie name. The child movie must contain `movieName` user data with the specified name.
- `kTargetChildMovieMovieID`. A child movie specified by the currently loaded movie's movie ID. The child movie must contain `movieID` user data with the specified ID.

Adding Wired Actions To a Flash Track

This section explains the steps you need to follow in order to add wired actions to a Macromedia Flash track. The Flash media handler was introduced in QuickTime 4 to enable a SWF 3.0 file to be treated as a track within a QuickTime movie. See “Flash Media” (page 167) for more information about the Flash media handler.

Sample code (`AddFlashActions`) is provided on the QuickTime SDK, as well as on the QuickTime developer website, that lets you add wired actions to a Flash track.

Note

For more detailed information about working with Flash, you can download the Macromedia SWF File Format Specification at <http://www.macromedia.com/software/flash/open/spec/>, along with the SWF File Parser code also at the Macromedia website. ♦

Extending the SWF Format

QuickTime 4 extends the SWF file format to allow the execution of any of its wired actions, in addition to the much smaller set of Flash actions. For example, you may use a SWF file as a user interface element in a QuickTime movie, controlling properties of the movie and other tracks. QuickTime also allows SWF files to be compressed using the zlib data compressor. This can significantly lower the bandwidth required when downloading a SWF file when it is in a QuickTime movie.

By using wired actions within a Flash track, compressing your Flash tracks, and combining Flash tracks with other types of QuickTime media, you can create compact and sophisticated multimedia content.

The SWF File Format Specification consists of a header followed by a series of tagged data blocks. The types of tagged data blocks you need to use are the `DefineButton2` and `DoAction`. The `DefineButton2` block allows Flash actions to be associated with a mouse state transition. `DoAction` allows actions to be executed when the tag is encountered. These are analogous to mouse-related QT event handlers and the frame loaded event in wired movies.

Some Useful Examples and Scenarios

Flash actions are stored in an action record. Each Flash action has its own tag, such as `ActionPlay` and `ActionNextFrame`. QuickTime defines one new tag: `QuickTimeActions`, which is `0xAA`. The data for the QuickTime actions tag is simply a QT atom container with the QuickTime wired actions to execute in it.

There are also fields you need to change in order to add wired actions to a SWF file. Additionally, there is one tag missing from the SWF file format that is described below.

What You Need to Modify

For `defineButton2`, you need to modify or add the following fields: file length, action records offset, the action offset, the condition, the record header size portion, and add action record.

File Length

A 32-bit field in the SWF file header.

`RecordHeader` for the `defineButton2`

`RecordHeader` contains the tag ID and length. You need to update the length. Note that there are short and long formats for record headers, depending on the size of the record. The tag ID for `defineButton2` is 34.

ActionRecordsOffset

The action records offset, a 16-bit field, is missing from the SWF File Format Specification. It occurs between the flags and buttons fields. It is initially set to 0 if there are no actions for the button. If there are actions for the button, then it must contain the offset from the point in the SWF file following this 16-bit value to the beginning of the action offset field.

```
DefineButton2 =
    Header
    ButtonID
    Flags
    ActionRecordsOffset    (this is missing from the spec)
    Buttons
    ButtonEndFlag
```

Some Useful Examples and Scenarios

```

Button2ActionCode
ActionOffset
Condition
Action          [ActionRecords]
ActionEndFlag

```

ActionOffset

There is one action offset per condition (mouse overDownToIdle). This is the offset used to skip over the condition and the following actions (the ActionRecord) for the condition. You need to update this value when adding actions.

Condition

The condition field is roughly equivalent to a wired movie event. The actions associated with button state transition condition are triggered when the transition occurs. You need to add or edit this field.

Actions

Flash actions each have their own action tag code. QuickTime actions use a single QuickTime actions code: 'AA'. You may add a list of actions to a single QuickTime actions tag.

The format of the QuickTime actions tag is as follows:

```

1 byte:      // Tag = 'AA'
2 bytes:     // data length (size of the QTAtomContainer)
n bytes      // the data which is the QTAtomContainer holding the
              // wired actions

```

DoAction

For DoAction, you need to modify a subset of the defineButton2 fields in the same manner as described above. These fields are file length, the record header size portion, and the action record.

Note that you need to write the length fields in little-endian format.

Creating Video Tracks at 30 Frames-per-Second

The duration of a video frame is stored in the time-to-sample atom contained within a sample table atom. This duration cannot be interpreted without the media's time scale, which defines the units-per-second for the duration. In this example, each frame has the same duration, so the time-to-sample atom has one entry, which applies to all video frames in the media.

As long as the ratio between frame duration and media time scale remains 1:30, any combination of values can be used for the duration and time scale. The larger the time scale the shorter the maximum duration. Since a movie defaults to a time scale of 600, this is a good number to use. It is also the least common multiple for 24, 25, and 30, making it handy for much of the math you are likely to encounter when making a movie.

The movie time scale is independent of the media time scale. Since you want to avoid movie edits that don't land on frame boundaries, it is a good idea to keep the movie time scale and the media time scale the same, or the movie time scale should be an even multiple of the media time scale. The movie time scale is stored in the movie header atom.

With a time scale of 600 in the media header atom, the time-to-sample atom would contain the following data values:

Atom size	24
Atom type	'stts'
Version / Flags	0
Number of entries	1
Sample count	n
Sample duration	20

Creating Video Tracks at 29.97 Frames-per-Second

NTSC color video is not 30 frames-per-second (fps), but actually 29.97 fps. The previous example showed how the media time scale and the duration of the

Some Useful Examples and Scenarios

frames specify the video's frame rate. By setting the media's time scale to 2997 units per second and setting the frame durations to 100 units each, the effective rate is 29.97 fps exactly.

In this situation, it is also a good idea to set the movie time scale to 2997 in order to avoid movie edits that don't land on frame boundaries. The movie's time scale is stored in the movie header atom.

With a time scale of 2997 in the media header atom, the time-to-sample atom would contain the following data values:

Atom size	24
Atom type	'stts'
Version/Flags	0
Number of entries	1
Sample count	n
Sample duration	100

Creating Audio Tracks at 44.1 Khz

The duration of an audio sample is stored in the time-to-sample atom contained in a sample table atom. This duration cannot be interpreted without the media's time scale, which defines the units-per-second for the duration. With audio, the duration of each audio sample is typically 1, so the time-to-sample atom has one entry, which applies to all audio samples.

With a time scale of 44100 in the media header atom, the time-to-sample atom would contain the following data values:

Atom size	24
Atom type	'stts'
Version/Flags	0
Number of entries	1
Sample count	n
Sample duration	1

This atom does not indicate whether the audio is stereo or mono or whether it contains 8-bit or 16-bit samples. That information is stored in the sound sample description atom, which is contained in the sample table atom.

Creating a Timecode Track for 29.97 FPS Video

A timecode track specifies timecode information for other tracks. The timecode keeps track of the timecodes of the original source of the video and audio. After a movie has been edited, the timecode can be extracted to determine the source tape and the timecodes of the frames.

It is important that the timecode track has the same time scale as the video track. Otherwise, the timecode will not tick at the exact same time as the video track.

For each contiguous source tape segment, there is a single timecode sample that specifies the timecode value corresponding to the start of the segment. From this sample, the timecode value can be determined for any point in the segment.

The sample description for a timecode track specifies the timecode system being used (for example, 30 fps drop frame) and the source information. Each sample is a timecode value.

Since the timecode media handler is a derived from the base media handler, the media information atom starts with a generic media header atom. The timecode atoms would contain the following data values:

Atom size	77
Atom type	'gmhd'
Atom size	69
Atom type	'gmin'
Version / Flags	0
Graphics mode	0x0040
Opcolor (red)	0x8000
Opcolor (green)	0x8000
Opcolor (blue)	0x8000
Balance	0

Some Useful Examples and Scenarios

Reserved	0
Atom size	45
Atom type	'tmcd'
Atom size	37
Atom type	'tcmi'
Version / Flags	0
Text font	0 (system font)
Text face	0 (plain)
Text size	12
Text color (red)	0
Text color (green)	0
Text color (blue)	0
Background color (red)	0
Background color (green)	0
Background color (blue)	0
Font name	'\pChicago' (Pascal string)

The sample table atom contains all the standard sample atoms and has the following data values:

Atom size	174
Atom type	'stbl' (sample table)
Atom size	74
Atom type	'std' (sample description)
Version / Flags	0
Number of entries	1
Sample description size [1]	58
Data format [1]	'tmcd'
Reserved [1]	0

Some Useful Examples and Scenarios

Data reference index [1]	1	
Flags[1]	0	
Flags (timecode) [1]	7 (drop frame + 24 hour + negative times OK)	
Time scale[1]	2997	
Frame duration[1]	100	
Number of frames [1]	20	
	Atom size	24
	Atom type	'name'
	String length	12
	Language code	0 (English)
	Name	"my tape name"
Atom size	24	
Atom type	'stts' (time to sample)	
Version/Flags	0	
Number of entries	1	
Sample count[1]	1	
Sample duration[1]	1	
Atom size	28	
Atom type	'stsc' (sample to chunk)	
Version/Flags	0	
Number of entries	1	
First chunk[1]	1	
Samples per chunk[1]	1	
Sample description ID[1]	1	
Atom size	20	
Atom type	'stsz' (sample size)	

Some Useful Examples and Scenarios

Version / Flags	0
Sample size	4
Number of entries	1
Atom size	20
Atom type	'stco' (chunk offset)
Version / Flags	0
Number of entries	1
Offset [1]	(offset into file of chunk 1)

In the example, let's assume that the segment's beginning timecode is 1:15:32.4 (1 hour, 15 minutes, 32 seconds, and 4 frames). The time would be expressed in the data file as 0x010F2004 (0x01 = 1 hour; 0x0F = 15 minutes; 0x20 = 32 seconds; 0x04 = 4 frames).

The video and audio tracks must contain a track reference atom to indicate that they reference this timecode track. The track reference is the same for both and is contained in the track atom (at the same level as the track header and media atoms).

This track reference would contain the following data values:

Atom size	12
Atom type	'tref'
Reference type	'tmcd'
Track ID of referenced track (timecode track)	3

In this example, the video and sound tracks are tracks 1 and 2. The timecode track is track 3.

Playing With Edit Lists

A segment of a movie can be repeated without duplicating media data by using edit lists. Suppose you have a single-track movie whose media time scale is 100 and track duration is 1000 (10 seconds). For this example the movie's time scale

Some Useful Examples and Scenarios

is 600. If there are no edits in the movie, the edit atom would contain the following data values:

Atom size	36
Atom type	'edts'
Atom size	28
Atom type	'elst'
Version / Flags	0
Number of entries	2
Track duration	6000 (10 seconds)
Media time	0
Media rate	1.0

Because this is a single-track move, the track's duration in the track header atom is 6000 and the movie's duration in the movie header atom is 6000.

If you change the track to play the media from time 0 to time 2 seconds, and then play the media from time 0 to time 10 seconds, the edit atom would now contain these data values:

Atom size	48
Atom type	'edts'
Atom size	40
Atom type	'elst'
Version / Flags	0
Number of entries	2
Track duration[1]	1200 (2 seconds)
Media time[1]	0
Media rate[1]	1.0
Track duration[2]	6000 (10 seconds)
Media time[2]	0
Media rate[2]	1.0

Because the track is now 2 seconds longer, the track's duration in the track header atom must now be 7200, and the movie's duration in the movie header atom must also be 7200.

Some Useful Examples and Scenarios

Currently, the media plays from time 0 to time 2, then plays from time 0 to time 10. If you take that repeated segment at the beginning (time 0 to time 2) and play it at double speed to maintain the original duration, the edit atom would now contain the following values:

Atom size	60
Atom type	'edts'
Atom size	52
Atom type	'elst'
Version / Flags	0
Number of entries	3
Track duration[1]	600 (1 second)
Media time[1]	0
Media rate[1]	2.0
Track duration[2]	600 (1 second)
Media time[2]	0
Media rate[2]	2.0
Track duration[3]	4800 (8 seconds)
Media time[3]	200
Media rate[3]	1.0

Because the track is now back to its original duration of 10 seconds, its duration in the track header atom is 6000, and the movie's duration in the movie header atom is 6000.


Interleaving Movie Data

In order to get optimal movie playback, you must create the movie with interleaved data. The data for the movie is placed on disk in time order so the data for a particular time in the movie is close together in the file. This means that you will have to intersperse the data from different tracks. To illustrate this, consider a movie with a single video and a single audio track.

Figure 5-7 shows how the movie data was collected, and how the data would need to be played back for proper synchronization. In this example, the video


data is recorded at 10 frames per second and the audio data is grouped into half-second chunks.

Figure 5-7 Non-interleaved movie data



After the data has been interleaved on the disk, the movie data atom would contain movie data in the order shown in Figure 5-8.

Figure 5-8 Interleaved movie data



In this example, the file begins with the movie atom ('moov'), followed by the movie data atom ('mdat'). In order to overcome any latencies in sound playback, at least one second of sound data is placed at the beginning of the interleaved data. This means that the sound and video data are offset from each other in the file by one second.

Referencing Two Data Files With a Single Track

The data reference index to be used for a given media sample is stored within that sample's sample description. Therefore, a track must contain multiple sample descriptions in order for that track to reference multiple data files. A different sample description must be used whenever the data file changes or whenever the format of the data changes. The sample-to-chunk atom determines which sample description to use for a sample.

The sample description atom would contain the following data values:

Atom size	...
Atom type	'stsd'
Version / Flags	0
Number of entries	2
Sample description size[1]	...
Data format	'tmcd'
Reserved	0
Data reference index	1
(sample data)	...
Sample description size[1]	...
Data format	'tmcd'
Reserved	0
Data reference index	2
(sample data)	...

If there is only 1 sample per chunk and the first 10 samples are extracted from sample description 2 and the next 30 samples are extracted from sample

Some Useful Examples and Scenarios

description 1, the sample-to-chunk atom would contain the following data values:

Atom size	40
Atom type	'stsc'
Version / Flags	0
Number of entries	2
First chunk[1]	1
Samples per chunk[1]	1
Sample description ID[1]	2
First chunk[2]	11
Samples per chunk[2]	1
Sample description ID[2]	1

The data reference atom would contain the following data values:

Atom size	...
Atom type	'dinf'
Atom size	...
Atom type	'dref'
Version / Flags	0
Number of entries	2
Size[1]	...
Type[1]	'alis'
Version[1]	0
Flags[1]	0 (not self referenced)
Data reference[1]	[alias pointing to file #1]
Size[2]	...
Type[2]	'rsrc'
Version[2]	0
Flags[2]	0 (not self referenced)
Data reference[2]	[alias pointing to file #2]

Getting the Name of a QuickTime VR Node

You can use standard QuickTime atom container functions to retrieve the information in a QuickTime VR node header atom. For example, the `MyGetNodeName` function defined in Listing 5-14 returns the name of a node, given its node ID.

Listing 5-14 Getting a node's name

```
OSErr MyGetNodeName (QTVRInstance theInstance, UInt32 theNodeID,
                                                             StringPtr theStringPtr)
{
    OSErr                theErr = noErr;
    QTAtomContainer      theNodeInfo;
    QTVRNodeHeaderAtomPtr theNodeHeader;
    QTAtom               theNodeHeaderAtom = 0;

    //Get the node information atom container.
    theErr = QTVRGetNodeInfo(theInstance, theNodeID, &theNodeInfo);

    //Get the node header atom.
    if (!theErr)
        theNodeHeaderAtom = QTFindChildByID(theNodeInfo, kParentAtomIsContainer,
                                              kQTVRNodeHeaderAtomType, 1, nil);
    if (theNodeHeaderAtom != 0) {
        QTLockContainer(theNodeInfo);

        //Get a pointer to the node header atom data.
        theErr = QTGetAtomDataPtr(theNodeInfo, theNodeHeaderAtom, nil,
                                  (Ptr *)&theNodeHeader);

        //See if there is a name atom.
        if (!theErr && theNodeHeader->nameAtomID != 0) {
            QTAtom theNameAtom;
            theNameAtom = QTFindChildByID(theNodeInfo, kParentAtomIsContainer,
                                          kQTVRStringAtomType, theNodeHeader->nameAtomID, nil);
            if (theNameAtom != 0) {
                VRStringAtomPtr theStringAtomPtr;
```

Some Useful Examples and Scenarios

```

//Get a pointer to the name atom data; copy it into the string.
theErr = QTGetAtomDataPtr(theNodeInfo, theNameAtom, nil,
                           (Ptr *)&theStringAtomPtr);

if (!theErr) {
    short theLen = theStringAtomPtr->stringLength;
    if (theLen > 255)
        theLen = 255;
    BlockMove(theStringAtomPtr->string, &theStringPtr[1], theLen);
    theStringPtr[0] = theLen;
}
}
}
}
QTUnlockContainer(theNodeInfo);
}

QTDisposeAtomContainer(theNodeInfo);
return(theErr);
}

```

The `MyGetNodeName` function defined in Listing 5-14 retrieves the node information atom container (by calling `QTVRGetNodeInfo`) and then looks inside that container for the node header atom with atom ID 1. If it finds one, it locks the container and then gets a pointer to the node header atom data. The desired information, the node name, is contained in the string atom whose atom ID is specified by the `nameAtomID` field of the node header structure. Accordingly, the `MyGetNodeName` function then calls `QTFindChildByID` once again to find that string atom. If the string atom is found, `MyGetNodeName` calls `QTGetAtomDataPtr` to get a pointer to the string atom data. Finally, `MyGetNodeName` copies the string data into the appropriate location and cleans up after itself before returning.

Adding Custom Atoms in a QuickTime VR Movie

If you author a QuickTime VR movie, you may choose to add custom atoms to either the VR world or node information atom containers. Those atoms can be extracted within an application to provide additional information that the application may use.

Some Useful Examples and Scenarios

Information that pertains to the entire scene might be stored in a custom atom within the VR world atom container. Node-specific information could be stored in the individual node information atom containers or as sibling atoms to the node location atoms within the VR world.

Custom hot spot atoms should be stored as siblings to the hot spot information atoms in the node information atom container. Generally, its atom type is the same as the custom hot spot type. You can set up an intercept procedure in your application in order to process clicks on the custom hot spots.

If you use custom atoms, you should install your hot spot intercept procedure when you open the movie. Listing 5-15 is an example of such an intercept procedure.

Listing 5-15 Typical hot spot intercept procedure

```
QTVRInterceptProc MyProc = NewQTVRInterceptProc (MyHotSpot);
QTVRInstallInterceptProc (qtvr, kQTVRTriggerHotSpotSelector, myProc, 0, 0);

pascal void MyHotSpot (QTVRInstance qtvr, QTVRInterceptPtr qtvrMsg,
                      Sint32 refCon, Boolean *cancel)
{
    UInt32 hotSpotID = (UInt32) qtvrMsg->parameter[0];
    QTAAtomContainer nodeInfo =
        (QTAAtomContainer) qtvrMsg->parameter[1];
    QTAAtom hotSpotAtom = (QTAAtom) qtvrMsg->parameter[2];
    OSType hotSpotType;
    CustomData myCustomData;
    QTAAtom myAtom;

    QTVRGetHotSpotType (qtvr, hotSpotID, &hotSpotType);
    if (hotSpotType != kMyAtomType) return;

    // It's our type of hot spot - don't let anyone else handle it
    *cancel = true;

    // Find our custom atom
    myAtom = QTFindChildByID (nodeInfo, hotSpotAtom, kMyAtomType, 1, nil);
    if (myAtom != 0) {
        OSErr err;
        // Copy the custom data into our structure
```

Some Useful Examples and Scenarios

```

err = QTCopyAtomDataToPtr (nodeInfo, myAtom, false,
                           sizeof(CustomData), &myCustomData, nil);
if (err == noErr)
    // Do something with it
    DoMyHotSpotStuff (hotSpotID, &myCustomData);
}
}

```

Your intercept procedure is called for clicks on any hot spot. You should check to see if it is your type of hot spot and if so, extract the custom hot spot atom and do whatever is appropriate for your hot spot type (`DoMyHotSpotStuff`).

When you no longer need the intercept procedure you should call `QTVRInstallInterceptProc` again with the same selector and a nil procedure pointer and then call `DisposeRoutineDescriptor` on `myProc`.

Apple reserves all hot spot and atom types with lowercase letters. Your custom hot spot type should contain all uppercase letters.

Adding Atom Containers in a QuickTime VR Movie

Assuming you have already created the QuickTime VR world and node information atom containers, Listing 5-16 shows the code (minus error checking) you would use to add them to the QTVR track.

Listing 5-16 Adding atom containers to a track

```

long descSize;
QTVRSampleDescriptionHandle qtvrSampleDesc;

// Create a QTVR sample description handle
descSize = sizeof(QTVRSampleDescription) + GetHandleSize((Handle) vrWorld) -
           sizeof(UInt32);
qtvrSampleDesc = (QTVRSampleDescriptionHandle) NewHandleClear (descSize);
(*qtvrSampleDesc)->size = descSize;
(*qtvrSampleDesc)->type = kQTVRQTVRType;

// Copy the VR world atom container data into the QTVR sample description

```

Some Useful Examples and Scenarios

```
BlockMove (*((Handle) vrWorld), &((*qtvrSampleDesc)->data),
          GetHandleSize((Handle) vrWorld));
// Now add it to the QTVR track's media
err = BeginMediaEdits (qtvrMedia);
err = AddMediaSample (qtvrMedia, (Handle) nodeInfo, 0,
    GetHandleSize((Handle) nodeInfo), duration,
    (SampleDescriptionHandle) qtvrSampleDesc, 1, 0, &sampleTime);
err = EndMediaEdits (qtvrMedia);
InsertMediaIntoTrack (qtvrTrack, trackTime, sampleTime, duration, 1L<<16);
```

The `duration` value is computed based on the duration of the corresponding image track samples for the node. The value of `trackTime` is the time for the beginning of the current node (zero for a single node movie). The values of `duration` and `sampleTime` are in the time base of the media; the value of `trackTime` is in the movie's time base.

Optimizing QuickTime VR Movies for Web Playback

Originally, both QuickTime movies and QuickTime VR movies had to be completely downloaded to the user's local hard disk before they could be viewed. Starting with QuickTime 2.5, if the movie data is properly laid out in the file, standard linear QuickTime movies can be viewed almost immediately. The frames that have been downloaded so far are shown while subsequent frames continue to be downloaded.

The important change that took place to allow this to happen was for QuickTime to place global movie information at the beginning of the file. Originally, it was at the end of the file. After that, the frame data simply needs to be in order in the file. Similarly, QuickTime VR files also need to be laid out in a certain manner in order to get some sort of quick feedback when viewing on the Web. Roughly speaking this involves writing out all of the media samples in the file in a particular order. Apple now provides a movie export component that does this for you: the QTVR Flattener.

The QTVR Flattener

The QTVR Flattener is a movie export component that converts an existing QuickTime VR single node movie into a new movie that is optimized for the

Some Useful Examples and Scenarios

Web. Not only does the flattener reorder the media samples, but for panoramas it also creates a small preview of the panorama. When viewed on the Web, this preview appears after 5% to 10% of the movie data has been downloaded, allowing users to see a lower-resolution version of the panorama.

Using the QTVR Flattener from your application is quite easy. After you have created the QuickTime VR movie, you simply open the QTVR Flattener component and call the `MovieExportToFile` routine as shown in Listing 5-17.

Listing 5-17 Using the flattener

```

ComponentDescription desc;
Component flattener;
ComponentInstance qtvrExport = nil;

desc.componentType = MovieExportType;
desc.componentSubType = MovieFileType;
desc.componentManufacturer = QTVRFlattenerType;

flattener = FindNextComponent(nil, &desc);
if (flattener) qtvrExport = OpenComponent (flattener);

if (qtvrExport)
    MovieExportToFile (qtvrExport, &myFileSpec, myQTVRMovie, nil, 0, 0);

```

The code fragment shown in Listing 5-17 creates a flattened movie file specified by the `myFileSpec` parameter. If your QuickTime VR movie is a panorama, the flattened movie file includes a quarter size, blurred JPEG, compressed preview of the panorama image.

Note

The constants `MovieExportType` and `MovieFileType` used in Listing 5-17 are defined in header files `QuickTimeComponents.h` and `Movies.h` respectively and are defined as 'spit' and 'MooV'. ♦

You can present users with the QTVR Flattener's own dialog box to allow them to choose options such as how to compress the preview image or to select a separate preview image file. Use the following code to show the dialog box:

```
err = MovieExportDoUserDialog (qtvrExport, myQTVRMovie, nil, 0, 0, &cancel);
```

Some Useful Examples and Scenarios

If the user cancels the dialog box, then the Boolean `cancel` is set to `true`.

If you do not want to present the user with the flattener's dialog box, you can communicate directly with the component by using the

`MovieExportSetSettingsFromAtomContainer` routine as described in the following paragraphs.

If you want to specify a preview image other than the default, you need to create a special atom container and then call

`MovieExportSetSettingsFromAtomContainer` before calling `MovieExportToFile`.

You can specify how to compress the image, what resolution to use, and you can even specify your own preview image file to be used. The atom container you pass in can have various atoms that specify certain export options. These atoms must all be children of a flattener settings parent atom.

The preview resolution atom is a 16-bit value that allows you to specify the resolution of the preview image. This value, which defaults to `kQTVRQuarterRes`, indicates how much to reduce the preview image.

The blur preview atom is a Boolean value that indicates whether to blur the image before compressing. Blurring usually results in a much more highly compressed image. The default value is `true`.

The create preview atom is a Boolean value that indicates whether a preview image should be created. The default value is `true`.

The import preview atom is a Boolean value that is used to indicate that the preview image should be imported from an external file rather than generated from the image in the panorama file itself. This allows you to have any image you want as the preview for the panorama. You can specify which file to use by also including the import specification atom, which is an `FSSpec` data structure that identifies the image file. If you do not include this atom, then the flattener presents the user with a dialog box asking the user to select a file. The default for import preview is `false`. If an import file is used, the image is used at its natural size and the resolution setting is ignored.

Sample Atom Container for the QTVR Flattener

The sample code in Listing 5-18 creates an atom container and adds atoms to indicate an import preview file for the flattener to use.

Listing 5-18 Specifying a preview file for the flattener to use

```

Boolean yes = true;
QTAtomContainer exportData;
QTAtom parent;
err = QTNewAtomContainer(&exportData);
// create a parent for the other settings atoms
err = QTInsertChild (exportData, kParentAtomIsContainer,
    QTVRFlattenerParentAtomType, 1, 0, 0, nil, &parent);
// Add child atom to indicate we want to import the preview from a file
err = QTInsertChild (exportData, parent, QTVRImportPreviewAtomType, 1, 0,
    sizeof (yes), &yes, nil);
// Add child atom to tell which file to import
err = QTInsertChild (exportData, parent, QTVRImportSpecAtomType, 1, 0,
    sizeof (previewSpec), &previewSpec, nil);
// Tell the export component
MovieExportSetSettingsFromAtomContainer (qtvExport, exportData);

```

Overriding the compression settings is a bit more complicated. You need to open a standard image compression dialog component and make calls to obtain an atom container that you can then pass to the QTVR Flattener component.

Listing 5-19 Overriding the compression settings

```

ComponentInstance sc;
QTAtomContainer compressorData;
SCSpatialSettings ss;
sc =
OpenDefaultComponent(StandardCompressionType, StandardCompressionSubType);
ss.codecType = kCinepakCodecType;
ss.codec = nil;
ss.depth = 0;
ss.spatialQuality = codecHighQuality
err = SCSetInfo(sc, scSpatialSettingsType, &ss);
err = SCGetSettingsAsAtomContainer(sc, &compressorData);
MovieExportSetSettingsFromAtomContainer (qtvExport, compressorData);

```


QuickTime Image File Format

This appendix describes QuickTime image files, which are intended to provide the most useful container for QuickTime-compressed still images.

Most still image file formats define both how images should be stored and compressed. However, the QuickTime image file format is a container format, which describes a storage mechanism independent of compression. The QuickTime image file format uses the same atom-based structure as a QuickTime movie.

Atom Types in QuickTime Image Files


There are two mandatory atom types: `'idsc'`, which contains an image description, and `'idat'`, which contains the image data. This is illustrated in Figure A-1. A QuickTime image file can also contain other atoms. For example, it can contain single-fork preview atoms.

In QuickTime 4, there is a new optional atom type: `'iicc'`, which can store a ColorSync profile.

QuickTime Image File Format

Table A-1 shows an example QuickTime image file containing a JPEG-compressed image.

Figure A-1 An 'idsc' atom followed by an 'idat' atom



QuickTime Image File Format

Table A-1 A QuickTime image file containing JPEG-compressed data

0000005E	Atom size, 94 bytes
69647363	Atom type, 'idsc'
00000056	Image description size, 86 bytes
6A706567	Compressor identifier, 'jpeg'
00000000	Reserved, set to 0
0000	Reserved, set to 0
0000	Reserved, set to 0
00000000	Major and minor version of this data, 0 if not applicable
6170706C	Vendor who compressed this data, 'appl'
00000000	Temporal quality, 0 (no temporal compression)
00000200	Spatial quality, <code>codecNormalQuality</code>
0140	Image width, 320
00F0	Image height, 240
00480000	Horizontal resolution, 72 dpi
00480000	Vertical resolution, 72 dpi
00003C57	Data size, 15447 bytes (use 0 if unknown)
0001	Frame count, 1
0C 50 68 6F 74 6F 20 2D 20 4A 50 45 47 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00	Compressor name, "Photo - JPEG" (32-byte Pascal string)
0018	Image bit depth, 24
FFFF	Color lookup table ID, -1 (none)
00003C5F	Atom size, 15455 bytes
69646174	Atom type, 'idat'
FF D8 FF E0 00 10 4A 46 49 46 00 01 01 01 00 48 ...	JPEG compressed data

IMPORTANT

The exact order and size of atoms is not guaranteed to

QuickTime Image File Format

match the example in Table A-1. Applications reading QuickTime image files should always use the atom size to traverse the file and ignore atoms of unrecognized types. ▲

Note

Like QuickTime movie files, QuickTime image files are big-endian. However, image data is stored in the same byte order as usually specified by the particular compression format. ◆

Recommended File Type and Suffix

Because the QuickTime image file is a single-fork format, it works well in cross-platform applications. On Mac OS systems, QuickTime image files are identified by the file type 'qtif'. Apple recommends using the filename extension .QIF to identify QuickTime image files on other platforms.

Defining Media Data Layouts

The QuickTime file format provides a great deal of flexibility in how media data is physically arranged within a file. However, it also allows media layouts to be created that may be inefficient for playback on a given device. To complicate the matter, a media layout that is inefficient for one device may be, in fact, very efficient for another. The purpose of this appendix is to define some common uses of QuickTime files and describe the media layout in these circumstances.

A QuickTime file can reference media data stored in a number of files, including the file itself. If a QuickTime file references only media data contained within itself, the file is said to be *self-contained*. A QuickTime file can also reference media data stored in files that are *not* QuickTime files. This is because the QuickTime file format references media within a URL by file offset, rather than by a data structuring mechanism of a particular file format. This allows a QuickTime file to refer to data stored in any container format.

Often, it is convenient to store a single media stream per file, for example, when encoding content. It is also useful for purposes of reusing content. (To reuse an elementary stream, it is not necessary to extract it from a larger, possibly multiplexed file.)

Because QuickTime can reference media stored in any file, it is not required that media be stored in the QuickTime file format. However, this is recommended. Putting the elementary streams in a QuickTime file has several advantages, particularly in enabling interchange of the content between different tools. Further, the QuickTime file format adds very little overhead to the media—as little as a few hundred bytes in many cases—so there is no great penalty in storage space.

One of the issues facing any device (a server or a local workstation) that is attempting to play back a QuickTime file in real time is the number of file seeks that must be performed.

It is possible to arrange the data in a QuickTime file to minimize, and potentially eliminate, any seeks during the course of normal playback. (Of course, random access and other kinds of interactivity require seeks.) Minimizing seeks is accomplished by interleaving the media data in the QuickTime file in such a way that the layout of the media in the file corresponds to the order in which the media data will be required. It is expected that most

Defining Media Data Layouts

servers, for example, will stream QuickTime media using the facilities of the hint tracks.

Take a scenario where the QuickTime file contains a single hint track that references an audio and a visual media stream. In order to eliminate all seeks, the hint track media must be interleaved with the audio and visual stream data. Because the hint track sample must always be read before the audio and visual media that it references, the hint track samples must always immediately precede the samples they reference.

A simple illustration of the ordering of data (that is, time and file offset increasing from left to right) is as follows:

```
H0 A0 H1 V1 H2 V2 H3 A1 H4 A2 V3 H5 V4
```

When a single hint sample references multiple pieces of media data, those pieces of media data must occur in the order that they are referenced.

Random Access

This appendix describes how to seek with a QuickTime file using child atoms.

Seeking with a QuickTime file is accomplished primarily by using the child atoms contained in the sample table atom. If an edit list is present, it must also be consulted. If you want to seek a given track to a time T , where T is in the time scale of the movie header atom, you could perform the following operations:

1. If the track contains an edit list, determine which edit contains the time T by iterating over the edits. The start time of the edit in the movie time scale must then be subtracted from the time T to generate T' , the duration into the edit in the movie time scale. T' is next converted to the time scale of the track's media to generate T'' . Finally, the time in the media scale to use is calculated by adding the media start time of the edit to T'' .
2. The time-to-sample atom for a track indicates what times are associated with which sample for that track. Use this atom to find the first sample prior to the given time.
3. The sample that was located in step 1 may not be a random access point. Locating the nearest random access point requires consulting two atoms. The sync sample table indicates which samples are in fact random access points. Using this table, you can locate which is the first sync sample prior to the specified time. The absence of the sync sample table indicates that all samples are synchronization points, and makes this problem easy. The shadow sync atom gives the opportunity for a content author to provide samples that are not delivered in the normal course of delivery, but which can be inserted to provide additional random access points. This improves random access without impacting bitrate during normal delivery. This atom maps samples that are not random access points to alternate samples which are. You should also consult this table if present to find the first shadow sync sample prior to the sample in question. Having consulted the sync sample table and the shadow sync table, you probably wish to seek to whichever resultant sample is closest to, but prior to, the sample found in step 1.
4. At this point you know the sample that will be used for random access. Use the sample-to-chunk table to determine in which chunk this sample is located.

APPENDIX C

Random Access

5. Knowing which chunk contained the sample in question, use the chunk offset atom to figure out where that chunk begins.
6. Starting from this offset, you can use the information contained in the sample-to-chunk atom and the sample size atom to figure out where within this chunk the sample in question is located. This is the desired information.

Meta-Data Handling

This appendix describes how meta-data is handled when QuickTime imports other file formats. (For more information about meta-data, refer to “Introduction to Atoms” (page 19) and “Compressed Movie Resources” (page 101)).

These formats are grouped into the following categories and sections:

- “Digital Video File Formats” (page 316)
- “Digital Audio File Formats” (page 316)
- “Still Image File Formats” (page 318)
- “Animation and 3D File Formats” (page 320)

Each section includes a table with specific details on the following, where applicable:

- the format supported by QuickTime—for example, the movie import component or the graphics import component
- the Macintosh file type—for example, 'Mp3 '
- file name extensions—for example, .mp3
- specific details for meta-data handling—for example, all Microsoft-defined “tombstone” data is transferred to the imported movie’s user data. Meta-data fields that have QuickTime equivalents are mapped as follows.
- software required—for example, QuickTime 3 or later

Digital Video File Formats

OpenDML and other AVI files	Description
Supported by	Movie import component
Macintosh file type	'VfW '
File name extensions	.avi
Meta-data handling	All Microsoft-defined “tombstone” data is transferred to the imported movie’s user data. Meta-data fields that have QuickTime equivalents are mapped as follows: 'ICOP' maps to kUserDataTextCopyright, 'ISBJ' maps to kUserDataTextInformation, 'INAM' maps to kUserDataTextFullName, 'ICRD' maps to '@day', 'IMED' maps to '@fmt', 'ISRC' maps to '@src'. Where no QuickTime equivalent exists, the metadata item’s four-character code is modified by replacing the initial I with a ©. All other characters remain unchanged.
Software required	QuickTime 3

Digital Audio File Formats

MPEG 1 layer 3	Description
Supported by	Movie import component
Macintosh file type	'Mp3 ', 'SwaT', 'MPEG', 'PLAY', 'MPG3', 'MP3 '

Meta-Data Handling

MPEG 1 layer 3

File name extensions

Meta-data handling

Software required

Description

.mp3, .swa

Metadata from ID3v1-style MP3 files is imported into the QuickTime movie.

Title maps to `kUserDataTextFullName`, artist maps to '@ART', album maps to '@alb', year maps to '@day', comment maps to '@cmt', and track number maps to '@des'.

QuickTime 4

WAV

Supported by

Macintosh file type

File name extensions

Meta-data handling

Software required

Description

Movie import component

'WAVE', '.WAV'

.wav

All Microsoft-defined “tombstone” data is transferred to the imported movie’s user data. Meta-data fields that have QuickTime equivalents are mapped as follows: 'ICOP' maps to `kUserDataTextCopyright`, 'ISBJ' maps to `kUserDataTextInformation`, 'INAM' maps to `kUserDataTextFullName`, 'ICRD' maps to '@day', 'IMED' maps to '@fmt', 'ISRC' maps to '@src'.

Where no QuickTime equivalent exists, the metadata item’s four-character code is modified by replacing the initial I with a @. All other characters remain unchanged.

QuickTime 2.5 or later

Still Image File Formats

FlashPix	Description
Supported by	Graphics import component
Macintosh file type	'FPix'
File name extensions	.fpx
Meta-data handling	Information about copyright, authorship, caption text, content description notes, camera manufacturer name, camera model name are transferred to kUserDataTextCopyright, kUserDataTextArtistField, kUserDataTextFullName, kParameterInfoWindowTitle, kParameterInfoManufacturer, kUserDataTextMakeField user data items, respectively.
Formats supported	1.0
Software required	QuickTime 4
GIF	Description
Supported by	Graphics import component
Macintosh file type	'GIFf', or 'GIF '
File name extensions	.gif
Meta-data handling	The GIF comment field is transferred to the kUserDataDateTextInformation user data item.
Software required	QuickTime 2.5 or later
JFIF/JPEG	Description
Supported by	Graphics import component
Macintosh file type	'JPEG'

Meta-Data Handling

JFIF/JPEG

File name extensions

Meta-data handling

Software required

Description

.jpg

The JFIF comment field is transferred to the imported Movie's user data in the `kUserDataTextInformation` field.

QuickTime 2.5 or later

Photoshop

Supported by

Macintosh file type

File name extensions

Meta-data handling

Description

Graphics import component

'8BPS'

.psd

Photoshop files store their metadata based on the IPTC-NAA Information Interchange Model and Digital Newsphoto Parameter Record. This information is transferred into the importer Movie's user data. The entire IPTC-NAA record is placed into a user data item of type 'iptc'. In addition, those metadata items which are defined by QuickTime are mapped directly to QuickTime types as follows: 116 to `kUserDataTextCopyright`, 120 to `kUserDataTextInformation`, 105 to `kUserDataTextFullName`, 55 to '@day', 115 to '@src'.

Software required

QuickTime 2.5 or later. QuickTime 3 is required for metadata handling.

QuickTime Image File

Supported by

Macintosh file type

File name extensions

Meta-data handling

Description

Graphics import component

'qtif'

.qtif, .qif, .qti

Metadata that is stored in `quickTimeImageFileMetaDataAtom` atom is copied directly to the Movie's user data.

Formats supported

All

Software required

QuickTime 2.5 or later

Meta-Data Handling

TIFF	Description
Supported by	Graphics Import Component
Macintosh file type	'TIFF'
File name extensions	.tif, .tiff
Meta-data handling	Extracted from standard tags and from IPTC block
Software required	QuickTime 3 or later

Animation and 3D File Formats

Animated GIF	Description
Supported by	Movie import component
Macintosh file type	'GIFf'
File name extensions	.gif
Meta-data handling	The GIF comment field is transferred to <code>kUserDataTextInformation</code> user data item.
Software required	QuickTime 3 or later

Summary of VR World and Node Atom Types

This appendix includes information that pertains to Chapter 3, “VR World Atom Container” (page 205) and “Node Information Atom Container” (page 211).

C Summary

Constants

VR World Atom Types

```
enum {
    kQTVRWorldHeaderAtomType    = FOUR_CHAR_CODE('vrsc'),
    kQTVRImagingParentAtomType  = FOUR_CHAR_CODE('imgp'),
    kQTVRPanoImagingAtomType    = FOUR_CHAR_CODE('impn'),
    kQTVRObjectImagingAtomType  = FOUR_CHAR_CODE('imob'),
    kQTVRNodeParentAtomType     = FOUR_CHAR_CODE('vrnp'),
    kQTVRNodeIDAtomType        = FOUR_CHAR_CODE('vrni'),
    kQTVRNodeLocationAtomType   = FOUR_CHAR_CODE('nloc')
};
```

Node Information Atom Types

```
enum {
    kQTVRNodeHeaderAtomType     = FOUR_CHAR_CODE('ndhd'),
    kQTVRHotSpotParentAtomType  = FOUR_CHAR_CODE('hspa'),
    kQTVRHotSpotAtomType       = FOUR_CHAR_CODE('hots'),
    kQTVRHotSpotInfoAtomType    = FOUR_CHAR_CODE('hsin'),
    kQTVRLinkInfoAtomType       = FOUR_CHAR_CODE('link')
};
```

Miscellaneous Atom Types

```
enum {
    kQTVRStringAtomType          = FOUR_CHAR_CODE('vrsg'),
    kQTVRPanoSampleDataAtomType  = FOUR_CHAR_CODE('pdat'),
    kQTVRObjectInfoAtomType      = FOUR_CHAR_CODE('obji'),
    kQTVRAltImageTrackRefAtomType = FOUR_CHAR_CODE('imtr'),
    kQTVRAltHotSpotTrackRefAtomType = FOUR_CHAR_CODE('hstr'),
    kQTVRAngleRangeAtomType      = FOUR_CHAR_CODE('arng'),
    kQTVRTrackRefArrayAtomType    = FOUR_CHAR_CODE('tref'),
    kQTVRPanConstraintAtomType    = FOUR_CHAR_CODE('pcon'),
    kQTVRTiltConstraintAtomType   = FOUR_CHAR_CODE('tcon'),
    kQTVRFOVConstraintAtomType    = FOUR_CHAR_CODE('fcon'),
    kQTVRCubicViewAtomType       = FOUR_CHAR_CODE('cuvw'),
    kQTVRCubicFaceDataAtomType    = FOUR_CHAR_CODE('cufa')
};
```

Track Reference Types

```
enum {
    kQTVRImageTrackRefType      = FOUR_CHAR_CODE('imgt'),
    kQTVRHotSpotTrackRefType    = FOUR_CHAR_CODE('hott')
};
```

Imaging Property Valid Flags

```
enum {
    kQTVRValidCorrection          = 1 << 0,
    kQTVRValidQuality            = 1 << 1,
    kQTVRValidDirectDraw         = 1 << 2,
    kQTVRValidFirstExtraProperty = 1 << 3
};
```

Link Hot Spot Valid Bits

```
enum {
    kQTVRValidPan                = 1 << 0,
    kQTVRValidTilt               = 1 << 1,
    kQTVRValidFOV                = 1 << 2,
    kQTVRValidViewCenter         = 1 << 3
};
```

Animation Settings

```
enum QTVRAnimationSettings {
    kQTVRObjectAnimateViewFramesOn          = (1 << 0),
    kQTVRObjectPalindromeViewFramesOn       = (1 << 1),
    kQTVRObjectStartFirstViewFrameOn        = (1 << 2),
    kQTVRObjectAnimateViewsOn               = (1 << 3),
    kQTVRObjectPalindromeViewsOn            = (1 << 4),
    kQTVRObjectSyncViewToFrameRate          = (1 << 5),
    kQTVRObjectDontLoopViewFramesOn         = (1 << 6),
    kQTVRObjectPlayEveryViewFrameOn        = (1 << 7)
};
```

Control Settings

```
enum QTVRControlSettings {
    kQTVRObjectWrapPanOn                    = (1 << 0),
    kQTVRObjectWrapTiltOn                   = (1 << 1),
    kQTVRObjectCanZoomOn                    = (1 << 2),
    kQTVRObjectReverseHControlOn            = (1 << 3),
    kQTVRObjectReverseVControlOn            = (1 << 4),
    kQTVRObjectSwapHVCControlOn             = (1 << 5),
    kQTVRObjectTranslationOn                = (1 << 6)
};
```

Controller Subtype and ID

```
enum {
    kQTControllerType                      = FOUR_CHAR_CODE('ctyp'),
    kQTControllerID                        = 1
};
```

Object Controller Types

```
enum ObjectUITypes {
    kGrabberScrollerUI                    = 1,
    kOldJoyStickUI                        = 2,
    kJoystickUI                           = 3,
    kGrabberUI                            = 4,
    kAbsoluteUI                           = 5
};
```

Node Location Flag

```
enum {
    kQTVRSameFile                = 0
};
```

Panorama Sample Flag

```
enum {
    kQTVRPanoFlagHorizontal      = 1 << 0,
    kQTVRPanoFlagAlwaysWrap      = 1 << 2
};
```

Data Types

```
typedef float                    Float32;
```

Sample Description Header Structure

```
typedef struct QTVRSampleDescription {
    UInt32                size;
    UInt32                type;
    UInt32                reserved1;
    UInt16                reserved2;
    UInt16                dataRefIndex;
    UInt32                data;
} QTVRSampleDescription, *QTVRSampleDescriptionPtr, **QTVRSampleDescriptionHandle;
```

String Atom Structure

```
typedef struct QTVRStringAtom {
    UInt16                stringUsage;
    UInt16                stringLength;
    unsigned char          string[4];
} QTVRStringAtom, *QTVRStringAtomPtr;
```

Summary of VR World and Node Atom Types

VR World Header Atom Structure

```
typedef struct QTVRWorldHeaderAtom {
    UInt16          majorVersion;
    UInt16          minorVersion;
    QAtomID         nameAtomID;
    UInt32          defaultNodeID;
    UInt32          vrWorldFlags;
    UInt32          reserved1;
    UInt32          reserved2;
} QTVRWorldHeaderAtom, *QTVRWorldHeaderAtomPtr;
```

Panorama-Imaging Atom Structure

```
typedef struct QTVRPanoImagingAtom {
    UInt16          majorVersion;
    UInt16          minorVersion;
    UInt32          imagingMode;
    UInt32          imagingValidFlags;
    UInt32          correction;
    UInt32          quality;
    UInt32          directDraw;
    UInt32          imagingProperties[6];
    UInt32          reserved1;
    UInt32          reserved2;
} QTVRPanoImagingAtom, *QTVRPanoImagingAtomPtr;
```

Node Location Atom Structure

```
typedef struct QTVRNodeLocationAtom {
    UInt16          majorVersion;
    UInt16          minorVersion;
    OSType          nodeType;
    UInt32          locationFlags;
    UInt32          locationData;
    UInt32          reserved1;
    UInt32          reserved2;
} QTVRNodeLocationAtom, *QTVRNodeLocationAtomPtr;
```

Summary of VR World and Node Atom Types

Node Header Atom Structure

```
typedef struct QTVRNodeHeaderAtom {
    UInt16          majorVersion;
    UInt16          minorVersion;
    OSType          nodeType;
    QTAtomID        nodeID;
    QTAtomID        nameAtomID;
    QTAtomID        commentAtomID;
    UInt32          reserved1;
    UInt32          reserved2;
} QTVRNodeHeaderAtom, *QTVRNodeHeaderAtomPtr;
```

Hot Spot Information Atom Structure

```
typedef struct QTVRHotSpotInfoAtom {
    UInt16          majorVersion;
    UInt16          minorVersion;
    OSType          hotSpotType;
    QTAtomID        nameAtomID;
    QTAtomID        commentAtomID;
    SInt32          cursorID[3];
    Float32         bestPan;
    Float32         bestTilt;
    Float32         bestFOV;
    FloatPoint      bestViewCenter;
    Rect            hotSpotRect;
    UInt32          flags;
    UInt32          reserved1;
    UInt32          reserved2;
} QTVRHotSpotInfoAtom, *QTVRHotSpotInfoAtomPtr;
```

Link Hot Spot Atom Structure

```
typedef struct QTVRLinkHotSpotAtom {
    UInt16          majorVersion;
    UInt16          minorVersion;
    UInt32          toNodeID;
    UInt32          fromValidFlags;
    Float32         fromPan;
    Float32         fromTilt;
```

Summary of VR World and Node Atom Types

```

Float32          fromFOV;
FloatPoint       fromViewCenter;
UInt32          toValidFlags;
Float32          toPan;
Float32          toTilt;
Float32          toFOV;
FloatPoint       toViewCenter;
Float32          distance;
UInt32          flags;
UInt32          reserved1;
UInt32          reserved2;
} QTVRLinkHotSpotAtom, *QTVRLinkHotSpotAtomPtr;

```

Angle Range Atom Structure

```

typedef struct QTVRAngleRangeAtom {
    Float32          minimumAngle;
    Float32          maximumAngle;
} QTVRAngleRangeAtom, *QTVRAngleRangeAtomPtr;

```

Panorama Sample Atom Structure

```

typedef struct QTVRPanoSampleAtom {
    UInt16          majorVersion;
    UInt16          minorVersion;
    UInt32          imageRefTrackIndex;
    UInt32          hotSpotRefTrackIndex;
    Float32          minPan;
    Float32          maxPan;
    Float32          minTilt;
    Float32          maxTilt;
    Float32          minFieldOfView;
    Float32          maxFieldOfView;
    Float32          defaultPan;
    Float32          defaultTilt;
    Float32          defaultFieldOfView;
    UInt32          imageSizeX;
    UInt32          imageSizeY;
    UInt16          imageNumFramesX;
    UInt16          imageNumFramesY;
    UInt32          hotSpotSizeX;

```

Summary of VR World and Node Atom Types

```

    UInt32          hotSpotSizeY;
    UInt16          hotSpotNumFramesX;
    UInt16          hotSpotNumFramesY;
    UInt32          flags;
    UInt32          reserved1;
    UInt32          reserved2;
} QTVRPanoSampleAtom, *QTVRPanoSampleAtomPtr;

```

Cubic View Atom Structure

```

struct QTVRCubicViewAtom {
    Float32          minPan;
    Float32          maxPan;
    Float32          minTilt;
    Float32          maxTilt;
    Float32          minFieldOfView;
    Float32          maxFieldOfView;

    Float32          defaultPan;
    Float32          defaultTilt;
    Float32          defaultFieldOfView;
};
typedef struct QTVRCubicViewAtom    QTVRCubicViewAtom;

```

Cubic Face Data Atom Structure

```

struct QTVRCubicFaceData {
    float    orientation[4];
    float    center[2];
    float    aspect;
    float    skew;
};
typedef struct QTVRCubicFaceData    QTVRCubicFaceData;

```

Object Sample Atom Structure

```

typedef struct QTVRObjectSampleAtom {
    UInt16          majorVersion;
    UInt16          minorVersion;
    UInt16          movieType;
    UInt16          viewStateCount;
}

```


APPENDIX E

Summary of VR World and Node Atom Types

```
    UInt16          defaultViewState;
    UInt16          mouseDownViewState;
    UInt32          viewDuration;
    UInt32          columns;
    UInt32          rows;
    Float32         mouseMotionScale;
    Float32         minPan;
    Float32         maxPan;
    Float32         defaultPan;
    Float32         minTilt;
    Float32         maxTilt;
    Float32         defaultTilt;
    Float32         minFieldOfView;
    Float32         fieldOfView;
    Float32         defaultFieldOfView;
    Float32         defaultViewCenterH;
    Float32         defaultViewCenterV;
    Float32         viewRate;
    Float32         frameRate;
    UInt32          animationSettings;
    UInt32          controlSettings;
} QTVRObjectSampleAtom, *QTVRObjectSampleAtomPtr;
```

Track Reference Entry Structure

```
struct QTVRTrackRefEntry {
    UInt32          trackRefType;
    UInt16          trackResolution;
    UInt32          trackRefIndex;
};
typedef struct QTVRTrackRefEntry QTVRTrackRefEntry;
```

A P P E N D I X E

Summary of VR World and Node Atom Types

Glossary

alpha channel The upper bits of a display pixel, which control the blending of video and graphical image data for a video digitizer component.

alternate track A movie **track** that contains alternate data for another track. QuickTime chooses one track to be used when the movie is played. The choice may be based on such considerations as image quality or localization.

API (Application Programming Interface) The set of function calls, data structures, and other programming elements by which a structure of code (such as a system-level toolbox) can be accessed by other code (such as an application program).

atom The basic unit of data in a **movie** resource, **sprite**, or other QuickTime data structure. There are a number of different atom types, including movie atoms, track atoms, and media atoms. There are two varieties of atoms: **QT atoms**, which may contain other atoms, and **classic atoms**, which do not contain any other atoms.

atom container A tree-structured hierarchy of **QT atoms**.

atom ID A 32-bit integer that uniquely identifies an **atom** among other **child atoms** of the same **parent atom**. The **root atom** has an atom ID value of 0x0001.

atom type A 32-bit value that uniquely identifies the data type of an atom. It is normally an `OSType`, rendered by four ASCII characters. An atom's data type helps determine how the atom's contents are interpreted.

background color The color of the background behind a sprite or other image.

bit depth The number of bits used to encode the color of each pixel in a graphics buffer.

chapter list A set of named entry points into a **movie**, presented to the viewer as a text list.

child atom A **QT atom** inside a **container atom**, which is its **parent atom**.

chunk A collection of sample data in a media. Chunks, which may contain one or more samples, allow optimized data access. Chunks in a media may have different sizes, and the samples within a chunk may have different sizes.

classic atom A QuickTime atom that contains no other atoms. A classic atom, however, may contain a table. An example of a classic atom is an edit list atom, containing the edit list table. Compare **QT atom**.

clipped movie boundary region The region that combines the union of all track movie boundary regions for a movie, which is the movie's **movie boundary region**,

with the movie's **movie clipping region**, which defines the portion of the movie boundary region that is to be used.

clipping The process of defining the boundaries of a graphics area.

container atom An **atom** that contains other atoms, possibly including other container atoms.

creator signature In the Macintosh file system, a four-character code that identifies the application program to which a file belongs.

data fork In a Macintosh file, the section that corresponds to a DOS/Windows file.

data handler A piece of software that is responsible for reading and writing a media's data. The data handler provides data input and output services to the media's **media handler**.

data reference A reference to a media's data.

display coordinate system The QuickDraw graphics world, which can be used to display QuickTime movies, as opposed to the movie's **time coordinate system**, which defines the basic time unit for each of the movie's tracks.

dithering A technique used to improve picture quality when you are attempting to display an image that exists at a higher bit-depth representation on a lower bit-depth device. For example, you might want to dither a 24 bits per pixel image for display on an 8-bit screen.

dropframe A synchronizing technique that skips timecodes to keep them current with video frames.

duration A time interval. Durations are time values that are interpreted as spans of time, rather than as points in time.

edit list A data structure that arranges a **media** into a time sequence.

edit state Information defining the current state of a movie or track with respect to an edit session. QuickTime uses edit states to support undo facilities.

effect description A data structure that specifies which component will be used to implement an effect in a movie, and how the component will be configured.

effect track A **modifier track** that applies an effect (such as a wipe or dissolve) to a movie.

file fork A section of a Macintosh file. See **data fork**, **resource fork**.

file preview A **thumbnail picture** from a movie that is displayed in the Open File dialog box.

fixed point A point that uses fixed-point numbers to represent its coordinates. QuickTime uses fixed points to provide greater display precision for graphical and image data.

fixed rectangle A rectangle that uses **fixed points** to represent its vertices. QuickTime uses fixed rectangles to provide greater display precision.

Flash A vector-based graphics and animation technology. Flash data is exported by **SWF files**.

flattening The process of copying all of the original data referred to by reference in QuickTime tracks into a QuickTime movie file. This can also be called *resolving references*. Flattening is used to bring in all of the data that may be referred to from multiple files after QuickTime editing is complete. It makes a QuickTime movie stand-alone—that is, it can be played on any system without requiring any additional QuickTime movie files or tracks, even if the original file referenced hundreds of files. The flattening operation is essential if QuickTime movies are to be used with CD-ROM discs.

frame A single image in a **sequence** of images.

frame rate The rate at which a movie is displayed—that is, the number of frames per second that are actually being displayed. In QuickTime the frame rate at which a movie was recorded may be different from the frame rate at which it is displayed. On very fast machines, the playback frame rate may be faster than the record frame rate; on slow machines, the playback frame rate may be slower than the record frame rate. Frame rates may be fractional.

free atom An **atom** of type 'free', which you can include in a QuickTime file as a placeholder for unused space.

graphics mode The method by which two overlapping images are blended together to produce a composite image.

graphics world A software environment in which a movie track or set of images may be defined before importing them into a movie.

handler reference atom A **QT atom** of type 'hdlr' that specifies the **media handler** to be used to interpret a **media**.

hint track A track in a **streaming movie** that contains information for a packetizer about the data units to stream.

hot spot An area, typically in a **VR** presentation, that the user can click to invoke an action.

hypertext Action media that contains a **URL** and takes the user to a website.

identity matrix A **transformation matrix** that specifies no change in the coordinates of the source image. The resulting image corresponds exactly to the source image.

image In **sprite** programming, one of a sprite's **properties**.

image sequence A series of visual representations usually represented by video over time. Image sequences may also be generated synthetically, such as from an animation sequence.

image track Any track in a **QuickTime** movie that contains visual images. The term particularly applies to video tracks that contain **VR** data.

input map A data structure that describes where to find information about tracks that are targets of a **modifier track**.

interlacing A video mode that updates half the scan lines on one pass and goes through the second half during the next pass.

interleaving A technique in which sound and video data are alternated in small pieces, so the data can be read off disk as it is needed. Interleaving allows for movies of almost any length with little delay on startup.

ISO Acronym for the International Standards Organization. ISO establishes standards for multimedia data formatting and transmission, such as **JPEG** and **MPEG**.

Joint Photographic Experts Group (JPEG) Refers to an international standard for compressing still images. This standard supplies the algorithm for image compression. The version of JPEG supplied with QuickTime complies with the baseline **ISO** standard bitstream, version 9R9. This algorithm is best suited for use with natural images.

key frame A sample in a sequence of temporally compressed samples that does not rely on other samples in the sequence for any of its information. Key frames are placed into temporally compressed sequences at a frequency that is determined by the **key frame rate**. Typically, the term *key frame* is used with respect to temporally compressed sequences of image data. See also **sync sample**.

key frame rate The frequency with which **key frames** are placed into temporally compressed data sequences.

leaf atom An **atom** that contains only data, and no other atoms.

layer A mechanism for prioritizing the tracks in a movie or the overlapping of sprites. When it plays a movie, QuickTime displays the movie's tracks according to

their layer—tracks with lower layer numbers are displayed first; tracks with higher layer numbers are displayed over those tracks.

matrix See **transformation matrix**.

matte A defined region of a movie display that can be clipped and filled with another display.

media A data structure that contains information that describes the data for a track in a movie. Note that a media does not contain its data; rather, a media contains a reference to its data, which may be stored on disk, CD-ROM disc, or any other mass storage device. Also called a *media structure*.

media handler A piece of software that is responsible for mapping from the movie's time coordinate system to the media's time coordinate system. The media handler also interprets the media's data. The **data handler** for the media is responsible for reading and writing the media's data.

MIDI Acronym for Musical Instrument Digital Interface, a standard format for sending instructions to a musical synthesizer.

modifier track A track in a movie that modifies the data or presentation of other tracks. For example, a **tween** track is a modifier track.

movie A structure of **time-based data** that is managed by QuickTime. A movie may contain sound, video, animation, or a combination of any of these types of data. A QuickTime movie contains one or more **tracks**; each track represents a single data stream in the movie.

movie boundary region A region that describes the area occupied by a movie in the movie coordinate system, before the movie has been clipped by the **movie clipping region**. A movie's boundary region is built up from the **track movie boundary regions** for each of the movie's **tracks**.

movie clipping region The clipping region of a movie in the movie's coordinate system. QuickTime applies the movie's clipping region to the **movie boundary region** to obtain a clipped movie boundary region. Only that portion of the movie that lies in the clipped movie boundary region is then transformed into an image in the display coordinate system.

movie display boundary region A region that describes the display area occupied by a movie in the display coordinate system, before the movie has been clipped by the **movie display clipping region**.

movie display clipping region The clipping region of a movie in the display coordinate system. Only that portion of the movie that lies in the clipping region is visible to the user. QuickTime applies the movie's display clipping region to the **movie display boundary region** to obtain the visible image.

movie file A QuickTime file that stores a movie and its associated data.

movie header atom A **QT atom** that specifies the characteristics of an entire QuickTime **movie**.

movie poster A single visual image representing a QuickTime movie. You specify a poster as a point in time in the movie and specify the tracks that are to be used to constitute the poster image.

movie preview A short dynamic representation of a QuickTime movie. Movie previews typically last no more than 3 to 5 seconds, and they should give the user some idea of what the movie contains. You define a movie preview by specifying its start time, its duration, and its tracks.

movie resource One of several data structures that provide the medium of exchange for movie data between applications on a Macintosh computer and between computers, even computers of different types.

movie sprite A **sprite** that lives in a **sprite track** and acts in a **movie**.

MPEG-4 A forthcoming **ISO** standard that is based on the QuickTime file format and will support video and audio **streaming**.

music One of the QuickTime media types, in which sequences of sounds and tones are generated.

National Television System Committee (NTSC) Refers to the color-encoding method adopted by the committee in 1953. This standard was the first monochrome-compatible, simultaneous color transmission system used for public broadcasting. This method is used widely in the United States.

node Either a panorama or an object in a QuickTime **VR movie**.

NTSC See **National Television System Committee**.

object track A track in a **QuickTime VR** movie that contains a set of views of a VR object.

offset-binary encoding A method of digitally encoding sound that represents the range of amplitude values as an unsigned number, with the midpoint of the range representing silence. For example, an 8-bit sound sample stored in offset-binary format would contain sample values ranging from 0 to 255, with a value of 128 specifying silence (no amplitude). Samples in Macintosh sound resources are stored in offset-binary form. Compare **twos-complement encoding**.

PAL See **Phase Alternation Line**.

panorama A structure of **QuickTime VR** data that forms a virtual-world environment within which the user can navigate.

panorama track A track in a **QuickTime VR** movie that contains a **panorama**.

parent atom A **QT atom** that contains other QT atoms, which are its **child atoms**.

Phase Alternation Line (PAL) A color-encoding system used widely in Europe, in which one of the subcarrier phases derived from the color burst is inverted in phase from one line to the next. This technique minimizes hue errors that may result during color video transmission. Sometimes called *Phase Alternating Line*.

playback quality A relative measure of the fidelity of a track in a QuickTime movie. You can control the playback (or language) quality of a movie during movie playback.

QuickTime chooses tracks from **alternate tracks** that most closely correspond to the display quality desired.

poster A frame shot from a **movie**, used to represent its content to the user.

preferred rate The default playback rate for a QuickTime movie.

preferred volume The default sound volume for a QuickTime movie.

preview A short, potentially dynamic, visual representation of the contents of a file. The Standard File Package can use previews in file dialog boxes to give the user a visual cue about a file's contents. See **file preview**.

preview atom An **atom** of type 'pnot', which can appear in a QuickTime file to contain a movie's **file preview**.

property Information about a **sprite** that describes its location or appearance. One sprite property is its **image**, the original bitmapped graphic of the sprite.

QT atom A QuickTime atom that contains other atoms, possibly including other QT and classic atoms. A data reference atom is an example of a QT atom. Compare **classic atom**.

QTMA (QuickTime Music Architecture) The part of QuickTime that lets other code create and manipulate music tracks in movies.

QTVR track A track in a QuickTime movie that maintains a list of **VR nodes**.

QuickDraw The original Mac OS two-dimensional drawing software, used by **QuickTime**.

QuickTime A set of Macintosh system extensions or a Windows dynamic-link library that other code can use to create and manipulate **time-based data**.

QuickTime VR A QuickTime media type that lets users interactively explore and examine photorealistic three-dimensional virtual worlds. QuickTime VR data structures are also called **panoramas**.

rate A value that specifies the pace at which time passes for a **time base**. A time base's rate is multiplied by the time scale to obtain the number of **time units** that pass per second. For example, consider a time base that operates in a time coordinate system that has a time scale of 60. If that time base has a rate of 1, 60 time units are processed per second. If the rate is set to 1/2, 30 time units pass per second. If the rate is 2, 120 time units pass per second.

resource In Macintosh programming, an entity in a file or in memory that may contain executable code or a description of a user interface item. Resources are loaded as needed by a resource manager, and are identified by their type and ID number.

resource fork In a Macintosh file, the section that contains **resources**.

root atom The largest **atom container** in a hierarchy, with **atom type** 'sean'.

sample A single element of a sequence of time-ordered data.

sample format The format of data samples in a track, such as a sprite track.

sample number A number that identifies the sample with data for a specified time.

SECAM (Systeme Electronique Couleur avec Memoire) Sequential Color With Memory; refers to a color-encoding system in which the red and blue color-difference information is transmitted on alternate lines, requiring a one-line memory in order to decode green information.

single-fork movie file A QuickTime movie file that stores both the movie data and the movie resource in the data fork of the movie file. You can use single-fork movie files to ease the exchange of QuickTime movie data between Macintosh computers and other computer systems. Compare **movie file**.

skip atom An **atom** of type 'skip', which you can include in a QuickTime file as a placeholder for unused space.

SMPTE Acronym for Society of Motion Picture and Television Engineers, an organization that sets video and movie technical standards.

sprite An animated image that is managed by QuickTime. A sprite is defined once and is then animated by commands that change its position or appearance.

sprite track A movie **track** populated by **movie sprites**.

streaming Delivery of video or audio data over a network in real time, to support applications such as videophone and video conferencing. See **MPEG-4**.

string atom An **atom** in **VR media** that contains text.

SWF files Files that contain **Flash** data.

sync sample A sample that does not rely on preceding frames for content. See also **key frame**.

Système Electronique Couleur avec Mémoire See **SECAM**.

temporal compression Image compression that is performed between **frames** in a sequence. This compression technique takes advantage of redundancy between adjacent frames in a sequence to reduce the amount of data that is required to accurately represent each frame in the sequence. Sequences that have been temporally compressed typically contain **key frames** at regular intervals. Compare **spatial compression**.

thumbnail picture A picture that can be created from an existing image that is stored as a pixel map, a picture, or a picture file. A thumbnail picture is useful for creating small representative images of a source image and in previews for files that contain image data.

time base A set of values that define the time basis for an entity, such as a QuickTime movie. A time base consists of a **time coordinate system** (that is, a **time scale** and a **duration**) along with a rate value. The rate value specifies the speed with which time passes for the time base.

time-based data Data that changes or interacts with the user along a time dimension. QuickTime is designed to handle time-based data.

timecode media A media of type 'tmcd' that is used to store timecode data.

timecode track A movie track that stores external timing information, such as **SMPTE** timecodes.

time coordinate system A set of values that defines the context for a **time base**. A time coordinate system consists of a **time scale** and a **duration**. Together, these values define the coordinate system in which a **time value** or a time base has meaning.

time scale The number of **time units** that pass per second in a **time coordinate system**. A time coordinate system that measures time in sixtieths of a second, for example, has a time scale of 60.

time unit The basic unit of measure for time in a time coordinate system. The value of the time unit for a time coordinate system is represented by the formula (1/time scale) seconds. A time coordinate system that has a time scale of 60 measures time in terms of sixtieths of a second.

time value A value that specifies a number of time units in a **time coordinate system**. A time value may contain information about a point in time or about a **duration**.

track A Movie Toolbox data structure that represents a single data stream in a QuickTime **movie**. A movie may contain one or more tracks. Each track is independent of other tracks in the movie and represents its own data stream. Each track has a corresponding **media**, which describes the data for the track.

track boundary region A region that describes the area occupied by a track in the track's coordinate system. QuickTime obtains this region by applying the **track**

clipping region and the **track matte** to the visual image contained in the **track rectangle**.

track clipping region The clipping region of a track in the track's coordinate system. QuickTime applies the track's clipping region and the **track matte** to the image contained in the **track rectangle** to obtain the **track boundary region**. Only that portion of the track that lies in the track boundary region is then transformed into an image in the movie coordinate system.

track header atom A **QT atom** that specifies the characteristics of a track in a QuickTime **movie**.

track height The height, in pixels, of the **track rectangle**.

track input map A structure of **QT atoms** that specifies how secondary data for a **track** is to be interpreted (clipping, blending, etc.).

track load settings Information that specifies how and when a **track** is to be preloaded before running in a **movie**.

track matte A pixel map that defines the blending of track visual data. The value of each pixel in the pixel map governs the relative intensity of the track data for the corresponding pixel in the result image. QuickTime applies the track matte, along with the **track clipping region**, to the image contained in the **track rectangle** to obtain the **track boundary region**.

track movie boundary region A region that describes the area occupied by a track in the movie coordinate system, before the movie has been clipped by the **movie clipping region**. The **movie boundary**

region is built up from the track movie boundary regions for each of the movie's **tracks**.

track offset The blank space that represents the intervening time between the beginning of a movie and the beginning of a track's data. In an audio track, the blank space translates to silence; in a video track, the blank space generates no visual image. All of the tracks in a movie use the movie's time coordinate system. That is, the movie's time scale defines the basic time unit for each of the movie's tracks. Each track begins at the beginning of the movie, but the track's data might not begin until some time value other than 0.

track reference A data structure that defines the relation between movie tracks, such as the relation between a **timecode track** and other tracks.

track rectangle A rectangle that completely encloses the visual representation of a track in a QuickTime movie. The width of this rectangle in pixels is referred to as the **track width**; the height, as the **track height**.

track width The width, in pixels, of the track rectangle.

transformation matrix A 3-by-3 matrix that defines how to map points from one coordinate space into another coordinate space.

tween data The data in a **tween track**, such as interpolation values.

tween track A **modifier track** that performs a specific kind of **tweening**, such as path-to-matrix rotation.

tweening A process interpolating new data between given values in conformance to an algorithm. It is an efficient way to expand or smooth a movie's presentation between its actual frames.

twos-complement encoding A system for digitally encoding sound that stores the amplitude values as a signed number—silence is represented by a sample with a value of 0. For example, with 8-bit sound samples, twos-complement values would range from -128 to 127, with 0 meaning silence. The Audio Interchange File Format (AIFF) stores samples in twos-complement form. Compare **offset-binary encoding**.

URL The address of a website.

user data Auxiliary data that your application can store in a QuickTime movie, track, or media structure. The user data is stored in a **user data list**; items in the list are referred to as **user data items**.

Examples of user data include a copyright, date of creation, name of a movie's director, and special hardware and software requirements.

user data item A single element in a **user data list**, such as a modification date or copyright notice.

user data list The collection of **user data** for a QuickTime movie, track, or media. Each element in the user data list is called a **user data item**.

VR (virtual reality) See **QuickTime VR**.

wired sprite A variety of sprite that acts like a button, telling other software when the user has clicked on its image or passed the cursor over it.

Bibliography

All the volumes of QuickTime API developer documentation are available online for download from Apple's QuickTime Technical Publications website at

<http://developer.apple.com/techpubs/quicktime/qtdevdocs/RM/AboutDocument.htm>

This website is the most current and up-to-date source for all QuickTime developer documentation. A complete roadmap of topics and functions is provided for developers who want to build applications using the QuickTime API. The QuickTime technical documentation suite consists of 12 books, totaling more than 7,000 pages.

QuickTime developer documents are also available in Adobe Portable Document format (PDF). PDF files can be opened and viewed online, as well as downloaded from

<http://developer.apple.com/techpubs/quicktime/qtdevdocs/RM/PDF.htm>

The QuickTime Developer Series

There are three books currently available in the QuickTime Developer Series. These books are published by Morgan Kaufmann.

Discovering QuickTime by George Towner (ISBN 0-12-059640-7), 515 pp. This is the official guide to programming QuickTime in C. It includes working demos and code samples on its CD-ROM.

QuickTime for Java by Tom Maremaa and William Stewart (ISBN 0-12-305440-0), 655 pp. This is the official guide for Java programmers who want to use QuickTime. It includes a full description of the Java interface to QuickTime and a CD-ROM with working examples.

QuickTime for the Web by Apple Computer (ISBN 0-12-471255-X), 720 pp. This is the most comprehensive book yet on authoring QuickTime movies on the Web.

Some Useful QuickTime Websites

Apple's official site for information, demos, sample code, online documentation, and the latest software:

<http://www.apple.com/quicktime/>

QuickTime terms and definitions current in the industry, plus a good set of links to other QuickTime sites:

<http://www.pcwebopedia.com/QuickTime.htm>

Channel QuickTime, an Apple home page with links to FAQs, forums, and the latest news about QuickTime:

<http://www.quicktimefaq.org/>

The entry point for a variety of announcement lists and discussion forums about QuickTime, multimedia, and other topics of interest to QuickTime developers:

<http://www.lists.apple.com/>

A site with lots of goodies, maintained by the authors of the MoviePlayer documentation and updated frequently. A useful resource:

<http://www.bmug.org/quicktime/>

The International QuickTime VR Association website, a professional association that promotes and supports the use of QuickTime VR and related technologies worldwide:

<http://www.iqtvra.org/>

QuickTime terms and definitions that are current in the industry, plus a good set of links to other QuickTime sites:

<http://www.pcwebopedia.com/QuickTime.htm>

Channel QuickTime, an Apple home page with links to FAQs, forums, and the latest news about QuickTime:

<http://www.quicktimefaq.org/>

Numerals

3D media 183

A

AddSpriteToSample function 269

AddTrackReference function 182, 279

'alis' data reference type 83

APP1 marker 119

'appl' data type 115

application markers 119

atom containers

creating 262

description key 156

disposing of 263

introduced 19

samples for QTVR Flattenner 304

structure 27

atom ID's 27

atoms

calculating sizes 24

child 27, 268

container 19, 27

creating 263

custom 210, 299, 301

IDs 27

introduction to 19

layout of 20

leaf 20, 27

parent 27

atom types

See also names of constants under K

'clik' 219

'clip' 38, 48

'cmov' 102

'cmdv' 102

'crsr' 211

'ctab' 38, 42

'CURS' 211

'dcom' 102

'dflt' 141

'dinf' 80, 192

'dref' 82, 192

'edts' 48, 56

'elst' 57, 294

'entr' 219

'exit' 219

'fiel' 117

'fram' 219

'free' 32

'from' 277

'gama' 117

'gmhd' 187

'hdlr' 67, 70

'hinf' 188

'hint' 190

'hnti' 188

'hots' 213

'hspa' 213

'htxt' 138

'idle' 219

'imag' 141

'imap' 63

'impn' 207

'kmat' 55

'link' 216

'matt' 48, 54

'mdat' 22, 32

'mdhd' 67, 69

'mdia' 48, 67

'minf' 67, 72, 187

'mjht' 117

'mjqt' 117

'moov' 32, 36, 38, 102

'mvhd' 38

'ndhd' 212

'nloc' 209

'pcnt' 276

'pnot' 32

'rtp' 188, 194

'sdp' 188

siSlopeAndIntercept 126

'skip' 32

'smhd' 72

'sprt' 141

'ssrc' 63

'stbl' 73, 85

- 'stco' 97
- 'stsc' 93
- 'std' 87, 191
- 'stss' 91
- 'stsz' 96
- 'stts' 88
- 'tcmi' 133
- 'tkhd' 48
- 'trak' 38, 48
- 'tref' 237
- 'trig' 219
- 'twdt' 169
- 'twdu' 169
- 'twen' 168
- 'twnt' 169
- 'twst' 169
- 'udta' 38, 67
- 'url' 218
- 'vmhd' 72
- 'vrcp' 211
- 'vrni' 209
- 'vrnp' 209
- 'vrsc' 206
- 'vrse' 204
- 'vrsg' 204
- 'wtxt' 138

- audio decompression 126
- audio tracks, creating 288

B

- balance values 259
- base media information header atoms 77
- big-endian format 310

C

- CBR audio. *See* constant bit-rate audio
- chapter lists 182
- 'chap' track reference type 183
- child atoms

- defined 27
- finding by index 268
- chunk offset atoms 97
- 'clik' atom type 219
- 'clip' atom type 38, 48
- clipping atoms 52
- clipping region atoms 53
- 'cmov' atom type 102
- 'cmvd' atom type 102
- color table atoms 42
- compressed matte atoms 55
- compressed movie resources 101
- constant bit-rate audio 127, 131
- container atoms 19, 27
- containment hierarchy 20
- 'crsr' atom type 211
- 'ctab' atom type 38, 42
- cubic panoramas 232
- 'CURS' atom type 211
- custom atoms
 - adding to QuickTime movies 299
 - for QTVR cursors 210
 - for QTVR hot spots 301
- 'cvid' compression type 115
- cylindrical panoramas 232

D

- data information atoms 80
- data reference atoms
 - examples of 297
 - fields in 82
- date and time values 256
- 'dcom' atom type 102
- DeleteTrackReference function 182
- 'dflt' atom type 141
- 'dinf' atom type 80, 192
- 'dref' atom type 82, 192

E

edit atoms 56
 edit list atoms 57, 292
 'edts' atom type 48, 56
 effect descriptions 273
 'elst' atom type 57, 294
 embedded movies 246, 282
 EMBED tag parameters 282
 'entr' atom type 219
 'exit' atom type 219
 external movie targets 281

F

'fiel' atom type 117
 Flash media 167, 284
 Flash media handler 284
 'fram' atom type 219
 'free' atom type 32
 free space atoms 32
 'from' atom type 277

G

'gama' atom type 117
 GetMovieBoundsRgn function 181
 'gmhd' atom type 187
 graphics display operations 256
 graphics mode
 blend graphics mode 258
 composition graphics mode 259
 copy graphics mode 258
 dither copy graphics mode 258
 premul black alpha graphics mode 259
 premul white alpha graphics mode 258
 straight alpha blend graphics mode 259
 straight alpha graphics mode 258
 transparent graphics mode 258
 graphics modes 257

H

handler reference atoms 70
 'hdlr' atom type 67, 70
 'hinf' atom type 188
 'hint' atom type 190
 hint media 185
 hint tracks
 and media tracks 190
 and sample description atoms 191
 and track atoms 48
 introduced 185
 RTP 191
 user data atoms 188
 'hnti' atom type 188
 'hots' atom type 213
 hot spot image tracks 236
 hot spot information atoms 214
 hot spot parent atoms 214
 hot spots in QTVR movies
 atom types 213
 sample code for intercepting 300
 'hspa' atom type 213
 'htxt' atom type 138
 Huffman tables 119
 hypertext for URLs 138

I

'idle' atom type 219
 'imag' atom type 141
 image compression formats 115
 image tracks
 defined 220
 low-resolution 236
 imaging parent atoms 207
 'imap' atom type 63
 'impn' atom type 207
 input maps
 creating 277
 updating 281
 interleaved movie data 294
 International QuickTime VR Association 342

interpolation tween tracks 173
ISO JPEG specification 119

J

JPEG 118
'jpeg' compression type 115

K

kActionTarget atom type 283
kCrossFadeTransitionType atom type 274
kDVAudioFormat sound codec 131
kEffectDataSourceType atom type 277
key frames 91, 100
kInitialRotationAtom atom type 176
kListElementType atom type 174
'kmat' atom type 55
kMovieMediaAltText atom type 249
kMovieMediaAutoPlay atom type 248
kMovieMediaBackgroundColor atom type 250
kMovieMediaClipBegin atom type 249
kMovieMediaClipDuration atom type 250
kMovieMediaDataReference atom type 246
kMovieMediaDefaultDataReferenceID atom type 246
kMovieMediaEnableFrameStepping atom type 250
kMovieMediaRectangleAtom atom type 251
kMovieMediaRegionAtom atom type 250
kMovieMediaSlaveAudio atom type 247
kMovieMediaSlaveGraphicsMode atom type 247
kMovieMediaSlaveTime atom type 246
kMovieMediaSlaveTrackDuration atom type 247
kMovieMediaTitle atom type 249
kMovieMediaUseMIMETYPE atom type 249
kNameAtom atom type 175
kParameterWhatName atom type 274
'kpcd' compression type 115
kQDesignCompression sound codec 131

kQTEventFrameLoaded atom type 219
kQTEventIdle atom type 219
kQTEventMouseClicked atom type 219
kQTEventMouseClickedEnd atom type 219
kQTEventMouseClickedEndTriggerButton atom type 219
kQTEventMouseEnter atom type 219
kQTEventMouseExit atom type 219
kQTVRColorCursorAtomType atom type 211
kQTVRCursorAtomType atom type 211
kQTVRCursorParentAtomType atom type 211
kQTVRHotSpotAtomType atom type 213
kQTVRHotSpotLinkType atom type 216
kQTVRHotSpotParentAtomType atom type 213
kQTVRHotSpotURLType atom type 218
kQTVRImagingParentAtomType atom type 207
kQTVRNodeHeaderAtomType atom type 212
kQTVRNodeIDAtomType atom type 209
kQTVRNodeLocationAtomType atom type 209
kQTVRNodeParentAtomType atom type 209
kQTVRPanoImagingAtomType atom type 207
kQTVRStringAtomType atom type 204
kQTVRTrackRefArrayAtomType atom type 237
kQTVRWorldHeaderAtomType atom type 206
kSpriteAtomType atom type 149
kSpriteBehaviorsAtomType atom type 149, 152
kSpriteCursorBehaviorAtomType atom type 149, 152
kSpriteFloatingPointVariableAtomType atom type 150
kSpriteImageAtomType atom type 149
kSpriteImageBehaviorAtomType atom type 149, 152
kSpriteImageDataAtomType atom type 149
kSpriteImageDataRefAtomType atom type 149, 153
kSpriteImageDataRefTypeAtomType atom type 149, 154
kSpriteImageDefaultImageIndexAtomType atom type 149, 154
kSpriteImageGroupIDAtomType atom type 149, 151, 153
kSpriteImageNameAtomType atom type 149
kSpriteImageRegistrationAtomType atom type 149, 151, 153

kSpriteImagesContainerAtomType **atom type** 149
 kSpriteNameAtomType **atom type** 149
 kSpriteSharedDataAtomType **atom type** 149, 150
 kSpriteStatusStringsBehaviorAtomType **atom type** 150, 152
 kSpriteStringVariableAtomType **atom type** 150
 kSpriteTrackPropertyHasActions **atom type** 148, 218
 kSpriteTrackPropertyQTIdleEventsFrequency **atom type** 148, 219
 kSpriteTrackPropertyScaleSpritesToScaleWorld **atom type** 149
 kSpriteTrackPropertyVisible **atom type** 149
 kSpriteURLLinkAtomType **atom type** 150
 kSpriteUsesImageIDsAtomType **atom type** 149, 152
 kSpriteVariablesContainerAtomType **atom type** 150
 kTargetParentMovie **atom type** 283
 kTargetRootMovie **atom type** 283
 kTrackModifierInput **atom type** 277
 kTrackModifierReference **atom type** 277, 279
 kTrackModifierType **atom type** 277
 kTween3dInitialCondition **atom type** 176
 kTweenData **atom type** 172
 kTweenDuration **atom type** 173, 174
 kTweenEntry **atom type** 172
 kTweenFlags **atom type** 173, 175
 kTweenInterpolationID **atom type** 173
 kTweenOutputMax **atom type** 173
 kTweenOutputMin **atom type** 173
 kTweenReturnDelta **atom type** 175
 kTweenStartOffset **atom type** 172
 kTweenType3dCameraData **atom type** 176
 kTweenType3dMatrix **atom type** 176
 kTweenType3dQuaternion **atom type** 176
 kTweenType3dRotateAboutAxis **atom type** 176
 kTweenType3dRotateAboutPoint **atom type** 176
 kTweenType3dRotateAboutVector **atom type** 176
 kTweenType3dRotatn **atom type** 176
 kTweenType3dScale **atom type** 176

kTweenType **atom type** 172, 175
 kTweenTypePathToFixedPoint **atom type** 175
 kTweenTypePathToMatrixRotation **atom type** 176
 kTweenTypePathToMatrixTranslationAndRotation **atom type** 175
 kTweenTypePathToMatrixTranslation **atom type** 175
 kTweenTypePathXtoY **atom type** 175
 kTweenTypePathYtoX **atom type** 175

L

language code values 253
 leaf atoms 20, 27
 'link' **atom type** 216
 link hot spot atoms 216
 low-resolution image tracks 236

M

MACE compression 129
 Macromedia Flash tracks 284
 matrices 256
 MatrixRecord structure 171
 'matt' **atom type** 48, 54
 'mdat' **atom type** 22, 32
 'mdhd' **atom type** 67, 69
 'mdia' **atom type** 48, 67
 media atoms 67
 media data atom types
 3D 183
 Flash 167
 hint 185
 modifier tracks 180
 movie 246
 MPEG 140
 music 139
 sound 124
 sprite 140
 streaming 184

- text 134
- timecode 131
- track references 181
- tween 168
- video media 114
- VR 203
- media header atoms 68
- media information atoms 72
- meta-data 101
- mime type for QuickTime files 31
- 'minf' atom type 67, 72, 187
- 'mjht' atom type 117
- 'mjpa' compression type 115
- 'mjpb' compression type 115
- 'mjqv' atom type 117
- modifier tracks
 - adding to movies 280
 - defined 180
- 'moov' atom type
 - as basic atom type 32, 36
 - as part of compressed movie 102
 - layout of 38
- 'MooV' file type 31
- Motion-JPEG
 - format A 119
 - format B 120
- .mov extension 31
- movie atoms
 - clipping 52
 - clipping region 53
 - color table 42
 - compressed matte 55
 - defined 35
 - layout of 38
 - movie header 39
 - track 47
 - track header 49
 - track matte 54
- movie data atoms 33
- movie header atoms 39
- movie media 246
- 'mpeg' compression type 115
- MPEG layer 3 codecs 131
- MPEG media 140
- 'MPEG' media type 140

- multinode movies 223
- music media
 - introduced 139
 - sample data 140
 - sample description 139
- 'musi' data type 139
- 'mvhd' atom type 38
- MyGetNodeName function 298

N

- 'ndhd' atom type 212
- 'nloc' atom type 209
- node information atom container 211
- node parent atoms 209
- nodes 223, 298

O

- object image tracks 244
- object tracks 220, 237, 244
- 'obje' data type 237

P

- panorama image tracks 221, 229
- panorama-imaging atoms 207
- panorama tracks 220, 225
- parent atoms 27
- 'pcnt' atom type 276
- 'plug' data type 282
- 'pnot' atom type 32
- preview atoms 33

Q

- 'qd3d' data type 183
- QT atom containers. *See* atom containers 27

QT atoms 24
QTCopyAtomDataToHandle function 270
QTCopyAtomDataToPtr function 270
QTCopyAtom function 267
QTCountChildrenOfType function 268
QTCreateStandardParameterDialog function 272
QTDisposeAtomContainer function 262
QTFindChildByID function 265
QTFindChildByIndex function 268
QTGetAtomTypeAndID function 270
QTGetNextChildType function 268
QTInsertChild function 263, 274, 279
QTInsertChildren function 265
QTLockContainer function 270
QTNextChildAnyType function 268
QTRemoveAtom function 271
QTRemoveChildren function 271
QTReplaceAtom function 267
QTSetAtomData function 270
QTSwapAtoms function 267
QTVRAngleRangeAtom structure 229
QTVR Flattener 302
'qtvr' media type 224
QTVRNodeHeaderAtom structure 212
QTVRNodeLocationAtom structure 209
QTVRObjectSampleAtom 238
QTVRPanoImagingAtom structure 208
QTVRTrackRefEntry structure 237
QTVR tracks 220
QTVRWorldHeaderAtom structure 206
QuickTime SDK 284
QuickTime VR file format 220
QuickTime VR movies
 adding atom containers to 302
 atoms in 203
 optimizing 303
 wired actions inn 218
QuickTime VR node header atoms 298

R

'raw' compression type 115

Real Time Transport Protocol 194
 See also RTP hint tracks
RGB data 118, 259
'rle' compression type 115
'rpza' compression type 115
'rsrc' data reference type 83
'rtp' atom type 188, 194
RTP hint tracks 191, 194

S

sample atoms
 introduced 83
 using 99
sample description atoms
 examples of 296
 introduced 87
 sound 124
 timecode 131
 video 114
samples
 and chunks 84
 defined 83
sample size atoms 96
sample table atoms 84
sample-to-chunk atoms
 defined 93
 examples of 297
'sdp' atom type 188
SetSpriteData function 268
siDecompressorSettings atom type 126
siSlopeAndIntercept atom type 126
'skip' atom type 32
'smc' compression type 115
'smhd' atom type 72
'soun' data type 72
sound codecs 127
SoundDescription structure 125
sound media 124
sound media handler 124
sound media information atoms 75
sound sample data 127
sound sample description 124

- sprite media
 - introduced 140
 - sample data 141
 - sample description 141
 - track format 145
 - tracks 143
- sprite track property atoms 148
- sprite tracks 143
 - 'sprt' atom type 141
 - 'sprt' media type 140
 - 'ssrc' atom type 63
 - 'stbl' atom type 73, 85
 - 'stco' atom type 97
- streaming media 184
- string atoms 204
- string encoding atom 204
 - 'stsc' atom type 93
 - 'stds' atom type 87, 191
 - 'stss' atom type 91
 - 'stsz' atom type 96
 - 'stts' atom type 88
 - 'svqi' compression type 115
- SWF file format 284
- sync sample atoms 91

T

- target atoms 282
 - 'tcmi' atom type 133
- text media
 - and hypertext 138
 - handler 134
 - introduced 134
 - sample data 137
 - sample description 134
- 'text' media type 134
- timecode media
 - handler 131
 - information atom 133
 - introduced 131
 - sample data 134
 - sample description 131
- timecode tracks, creating 289

- time-to-sample atoms 88
- time values 256
 - 'tkhd' atom type 48
 - 'tmcd' data type 131
- track atoms 47
- track header atoms 49
- track input map atoms 63
- track matte atoms 54
- track reference atoms 277, 292
- track references 181
 - 'trak' atom type 38, 48
- transfer modes. *See* graphics modes
- transformation matrices 256
 - 'tref' atom type 237
 - 'trig' atom type 219
 - 'twdt' atom type 169
 - 'twdu' atom type 169
- tween atoms
 - 3D 176
 - interpolation 177
 - introduced 172
 - list 176
 - path 175
 - region 178
 - sequence 179
- tween media
 - data 168
 - introduced 168
 - operation 168
 - QT atom container 172
 - sample data 168
 - sample description 168
 - track 168
 - type categories 171
 - type characteristics 170
 - type values 169
- tween tracks
 - defined 168
 - modifying 173
- 'twen' atom type 168
- 'twnt' atom type 169
- two-dimensional transformations 256
 - 'twst' atom type 169

U

- 'udta' atom type 38, 67
- 'ulaw' compression 128
- uncompressed audio 127
- uncompressed RGB data 118
- uncompressed yuv2 data 118
- 'url' atom type 218
- 'url' data reference type 83
- URL hot spot atoms 218
- user data list entry types 45

V

- VBR (variable bit-rate) audio 127, 131
- 'vide' data type 72, 114
- video media
 - data 114
 - handler 115
 - introduced 114
 - sample description 114
- video media information atoms 73
- video media information header atoms 74
- "video/quicktime" mime type 31
- video sample data
 - introduced 117
 - JPEG 118
 - Motion-JPEG 119
 - uncompressed RGB 118
 - uncompressed yuv2 118
- video sample descriptions
 - extending 116
 - introduced 114
- video tracks, creating 287
- 'vmhd' atom type 72
- 'vrcp' atom type 211
- VRLinkHotSpotAtom structure 216
- VR media 203
 - 'vrni' atom type 209
 - 'vrnp' atom type 209
 - 'vrsc' atom type 206
 - 'vrse' atom type 204
 - 'vrsg' atom type 204

- VR world atom container 205

W

- wired action grammar 159
- wired actions in QTVR tracks 218
- 'wtxt' atom type 138

Y

- 'Yuv2' compression type 115
- yuv2 video format 118